

# Fixed-delay Events in Generalized Semi-Markov Processes Revisited<sup>★</sup>

Tomáš Brázdil, Jan Krčál, Jan Křetínský<sup>★★</sup>, and Vojtěch Řehák

Faculty of Informatics, Masaryk University, Brno, Czech Republic  
{brazdil, krcal, jan.kretinsky, rehak}@fi.muni.cz

**Abstract.** We study long run average behavior of generalized semi-Markov processes with both fixed-delay events as well as variable-delay events. We show that allowing two fixed-delay events and one variable-delay event may cause an unstable behavior of a GSMP. In particular, we show that a frequency of a given state may not be defined for almost all runs (or more generally, an invariant measure may not exist). We use this observation to disprove several results from literature. Next we study GSMP with at most one fixed-delay event combined with an arbitrary number of variable-delay events. We prove that such a GSMP always possesses an invariant measure which means that the frequencies of states are always well defined and we provide algorithms for approximation of these frequencies. Additionally, we show that the positive results remain valid even if we allow an arbitrary number of reasonably restricted fixed-delay events.

## 1 Introduction

Generalized semi-Markov processes (GSMP), introduced by Matthes in [22], are a standard model for discrete-event stochastic systems. Such a system operates in continuous time and reacts, by changing its state, to occurrences of events. Each event is assigned a random delay after which it occurs; state transitions may be randomized as well. Whenever the system reacts to an event, new events may be scheduled and pending events may be discarded. To get some intuition, imagine a simple communication model in which a server sends messages to several clients asking them to reply. The reaction of each client may be randomly delayed, e.g., due to latency of communication links. Whenever a reply comes from a client, the server changes its state (e.g., by updating its database of alive clients or by sending another message to the client) and then waits for the rest of the replies. Such a model is usually extended by allowing the server to time-out and to take an appropriate action, e.g., demand replies from the remaining clients in a more urgent way. The time-out can be seen as another event which has a fixed delay.

More formally, a GSMP consists of a set  $S$  of states and a set  $\mathcal{E}$  of events. Each state  $s$  is assigned a set  $\mathbf{E}(s)$  of events *scheduled* in  $s$ . Intuitively, each event in  $\mathbf{E}(s)$  is assigned a positive real number representing the amount of time which elapses before

---

<sup>★</sup> The authors are supported by the Institute for Theoretical Computer Science, project No. 1M0545, the Czech Science Foundation, grant No. P202/10/1469 (T. Brázdil, V. Řehák) and No. 102/09/H042 (J. Krčál), and Brno PhD Talent Financial Aid (J. Křetínský).

<sup>★★</sup> On leave at TU München, Boltzmannstr. 3, Garching, Germany.

the event occurs. Note that several events may occur at the same time. Once a set of events  $E \subseteq \mathbf{E}(s)$  occurs, the system makes a *transition* to a new state  $s'$ . The state  $s'$  is randomly chosen according to a fixed distribution which depends only on the state  $s$  and the set  $E$ . In  $s'$ , the *old* events of  $\mathbf{E}(s) \setminus \mathbf{E}(s')$  are discarded, each *inherited* event of  $(\mathbf{E}(s') \cap \mathbf{E}(s)) \setminus E$  remains scheduled to the same point in the future, and each *new* event of  $(\mathbf{E}(s') \setminus \mathbf{E}(s)) \cup (\mathbf{E}(s') \cap E)$  is newly scheduled according to its given probability distribution.

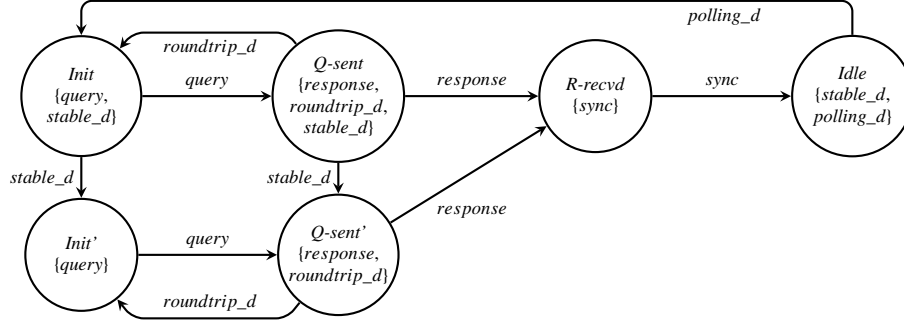
In order to deal with GSMP in a rigorous way, one has to impose some restrictions on the distributions of delays. Standard mathematical literature, such as [14,15], usually considers GSMP with continuously distributed delays. This is certainly a limitation, as some systems with fixed time delays (such as time-outs or processor ticks) cannot be faithfully modeled using only continuously distributed delays. We show some examples where fixed delays exhibit qualitatively different behavior than any continuously distributed approximation. In this paper we consider the following two types of events:

- *variable-delay*: the delay of the event is randomly distributed according to a probability density function which is continuous and positive either on a bounded interval  $[\ell, u]$  or on an unbounded interval  $[\ell, \infty)$ ;
- *fixed-delay*: the delay is set to a fixed value with probability one.

The desired behavior of systems modeled using GSMP can be specified by various means. One is often interested in long-run behavior such as mean response time, frequency of errors, etc. (see, e.g., [1]). For example, in the above communication model, one may be interested in average response time of clients or in average time in which all clients eventually reply. Several model independent formalisms have been devised for expressing such properties of continuous time systems. For example, a well known temporal logic CSL contains a steady state operator expressing frequency of states satisfying a given subformula. In [9], we proposed to specify long-run behavior of a continuous-time process using a timed automaton which observes runs of the process, and measure the frequency of locations of the automaton.

In this paper we consider a standard performance measure, the frequency of states of the GSMP. To be more specific, let us fix a state  $\hat{s} \in S$ . We define a random variable  $\mathbf{d}$  which to every run assigns the (discrete) frequency of visits to  $\hat{s}$  on the run, i.e. the ratio of the number of transitions entering  $\hat{s}$  to the number of all transitions. We also define a random variable  $\mathbf{c}$  which gives timed frequency of  $\hat{s}$ , i.e. the ratio of the amount of time spent in  $\hat{s}$  to the amount of time spent in all states. Technically, both variables  $\mathbf{d}$  and  $\mathbf{c}$  are defined as limits of the corresponding ratios on prefixes of the run that are prolonged ad infinitum. Note that the limits may not be defined for some runs. For example, consider a run which alternates between  $\hat{s}$  and another state  $s$ ; it spends 2 time unit in  $\hat{s}$ , then 4 in  $s$ , then 8 in  $\hat{s}$ , then 16 in  $s$ , etc. Such a run does not have a limit ratio between time spent in  $\hat{s}$  and in  $s$ . We say that  $\mathbf{d}$  (or  $\mathbf{c}$ ) is well-defined for a run if the limit ratios exist for this run. Our goal is to characterize stable systems that have the variables  $\mathbf{d}$  and  $\mathbf{c}$  well-defined for almost all runs, and to analyze the probability distributions of  $\mathbf{d}$  and  $\mathbf{c}$  on these stable systems.

As a working example of GSMP with fixed-delay events, we present a simplified protocol for time synchronization. Using the variable  $\mathbf{c}$ , we show how to measure reliability of the protocol. Via message exchange, the protocol sets and keeps a client clock



**Fig. 1.** A GSMP model of a clock synchronization protocol. Below each state label, we list the set of scheduled events. We only display transitions that can take place with non-zero probability.

sufficiently close to a server clock. Each message exchange is initialized by the client asking the server for the current time, i.e. sending a *query* message. The server adds a timestamp into the message and sends it back as a *response*. This query-response exchange provides a reliable data for *synchronization* action if it is realized within a given *round-trip delay*. Otherwise, the client has to repeat the procedure. After a success, the client is considered to be synchronized until a given *stable-time delay* elapses. Since the aim is to keep the clocks synchronized all the time, the client restarts the synchronization process sooner, i.e. after a given *polling delay* that is shorter than the stable-time delay. Notice that the client gets desynchronized whenever several unsuccessful synchronizations occur in a row. Our goal is to measure the portion of the time when the client clock is not synchronized.

Figure 1 shows a GSMP model of this protocol. The delays specified in the protocol are modeled using fixed-delay events *roundtrip\_d*, *stable\_d*, and *polling\_d* while actions are modeled by variable-delay events *query*, *response*, and *sync*. Note that if the stable-time runs out before a fast enough response arrives, the systems moves into primed states denoting it is not synchronized at the moment. Thus,  $\mathbf{c}(\text{Init}') + \mathbf{c}(\text{Q-sent}')$  expresses the portion of the time when the client clock is not synchronized.

**Our contribution.** So far, GSMP were mostly studied with variable-delay events only. There are a few exceptions such as [4,3,8,2] but they often contain erroneous statements due to presence of fixed-delay events. Our goal is to study the effect of mixing a number of fixed-delay events with an arbitrary amount of variable-delay events.

At the beginning we give an example of a GSMP with two fixed-delay events for which it is *not true* that the variables  $\mathbf{d}$  and  $\mathbf{c}$  are well-defined for almost all runs. We also disprove some crucial statements of [3,4]. In particular, we show an example of a GSMP which reaches one of its states with probability less than one even though the algorithms of [3,4] return the probability one. The mistake of these algorithms is fundamental as they neglect the possibility of unstable behavior of GSMP.

Concerning positive results, we show that if there is at most one fixed-delay event, then both  $\mathbf{d}$  and  $\mathbf{c}$  are almost surely well-defined. This is true even if we allow an arbitrary number of reasonably restricted fixed-delay events. We also show how to approxi-

mate distribution functions of  $\mathbf{d}$  and  $\mathbf{c}$ . To be more specific, we show that for GSMP with at most one unrestricted and an arbitrary number of restricted fixed-delay events, both variables  $\mathbf{d}$  and  $\mathbf{c}$  have finite ranges  $\{d_1, \dots, d_n\}$  and  $\{c_1, \dots, c_n\}$ . Moreover, all values  $d_i$  and  $c_i$  and probabilities  $\mathcal{P}(\mathbf{d} = d_i)$  and  $\mathcal{P}(\mathbf{c} = c_i)$  can be effectively approximated.

**Related work.** There are two main approaches to the analysis of GSMP. One is to restrict the amount of events or types of their distributions and to solve the problems using symbolic methods [8,2,20]. The other is to estimate the values of interest using simulation [26,14,15]. Concerning the first approach, time-bounded reachability has been studied in [2] where the authors restricted the delays of events to so called expolynomial distributions. The same authors also studied reachability probabilities of GSMP where in each transition at most one event is inherited [8]. Further, the widely studied formalisms of semi-Markov processes (see, e.g., [19,9]) and continuous-time Markov chains (see, e.g., [6,7]) are both subclasses of GSMP.

As for the second approach, GSMP are studied by mathematicians as a standard model for discrete event simulation and Markov chains Monte Carlo (see, e.g., [13,16,24]). Our work is strongly related to [14,15] where the long-run average behavior of GSMP with variable-delay events is studied. Under relatively standard assumptions the stochastic process generated by a GSMP is shown to be irreducible and to possess an invariant measure. In such a case, the variables  $\mathbf{d}$  and  $\mathbf{c}$  are almost surely constant. Beside the theoretical results, there exist tools that employ simulation for model checking (see, e.g., [26,10]).

In addition, GSMP are a proper subset of stochastic automata, a model of concurrent systems (see, e.g., [11]). Further, as shown in [15], GSMP have the same modeling power as stochastic Petri nets [21]. The formalism of deterministic and stochastic Petri nets (DSPN) introduced by [20] adds deterministic transitions – a counterpart of fixed-delay events. The authors restricted the model to at most one deterministic transition enabled at a time and to exponentially distributed timed transitions. For this restricted model, the authors proved existence of a steady state distribution and provided an algorithm for its computation. However, the methods inherently rely on the properties of the exponential distribution and cannot be extended to our setting with general variable delays. DSPN have been extended by [12,18] to allow arbitrarily many deterministic transitions. The authors provide algorithms for steady-state analysis of DSPN that were implemented in the tool DSPNExpress [17], but do not discuss under which conditions the steady-state distributions exist.

## 2 Preliminaries

In this paper, the sets of all positive integers, non-negative integers, real numbers, positive real numbers, and non-negative real numbers are denoted by  $\mathbb{N}$ ,  $\mathbb{N}_0$ ,  $\mathbb{R}$ ,  $\mathbb{R}_{>0}$ , and  $\mathbb{R}_{\geq 0}$ , respectively. For a real number  $r \in \mathbb{R}$ ,  $\text{int}(r)$  denotes its integral part, i.e. the largest integer smaller than  $r$ , and  $\text{frac}(r)$  denotes its fractional part, i.e.  $r - \text{int}(r)$ . Let  $A$  be a finite or countably infinite set. A *probability distribution* on  $A$  is a function  $f : A \rightarrow \mathbb{R}_{\geq 0}$  such that  $\sum_{a \in A} f(a) = 1$ . The set of all distributions on  $A$  is denoted by  $\mathcal{D}(A)$ .

A  $\sigma$ -field over a set  $\Omega$  is a set  $\mathcal{F} \subseteq 2^\Omega$  that includes  $\Omega$  and is closed under complement and countable union. A *measurable space* is a pair  $(\Omega, \mathcal{F})$  where  $\Omega$  is a set called *sample space* and  $\mathcal{F}$  is a  $\sigma$ -field over  $\Omega$  whose elements are called *measurable sets*. Given a measurable space  $(\Omega, \mathcal{F})$ , we say that a function  $f : \Omega \rightarrow \mathbb{R}$  is a random variable if the inverse image of any real interval is a measurable set. A *probability measure* over a measurable space  $(\Omega, \mathcal{F})$  is a function  $\mathcal{P} : \mathcal{F} \rightarrow \mathbb{R}_{\geq 0}$  such that, for each countable collection  $\{X_i\}_{i \in I}$  of pairwise disjoint elements of  $\mathcal{F}$ , we have  $\mathcal{P}(\bigcup_{i \in I} X_i) = \sum_{i \in I} \mathcal{P}(X_i)$  and, moreover,  $\mathcal{P}(\Omega) = 1$ . A *probability space* is a triple  $(\Omega, \mathcal{F}, \mathcal{P})$ , where  $(\Omega, \mathcal{F})$  is a measurable space and  $\mathcal{P}$  is a probability measure over  $(\Omega, \mathcal{F})$ . We say that a property  $A \subseteq \Omega$  holds for *almost all* elements of a measurable set  $Y$  if  $\mathcal{P}(Y) > 0$ ,  $A \cap Y \in \mathcal{F}$ , and  $\mathcal{P}(A \cap Y \mid Y) = 1$ . Alternatively, we say that  $A$  holds *almost surely* for  $Y$ .

## 2.1 Generalized semi-Markov processes

Let  $\mathcal{E}$  be a finite set of *events*. To every  $e \in \mathcal{E}$  we associate the *lower bound*  $\ell_e \in \mathbb{N}_0$  and the *upper bound*  $u_e \in \mathbb{N} \cup \{\infty\}$  of its delay. We say that  $e$  is a *fixed-delay* event if  $\ell_e = u_e$ , and a *variable-delay* event if  $\ell_e < u_e$ . Furthermore, we say that a variable-delay event  $e$  is *bounded* if  $u_e \neq \infty$ , and *unbounded*, otherwise. To each variable-delay event  $e$  we assign a *density function*  $f_e : \mathbb{R} \rightarrow \mathbb{R}$  such that  $\int_{\ell_e}^{u_e} f_e(x) dx = 1$ . We assume  $f_e$  to be positive and continuous on the whole  $[\ell_e, u_e]$  or  $[\ell_e, \infty)$  if  $e$  is bounded or unbounded, respectively, and zero elsewhere. We require that  $f_e$  have finite expected value, i.e.  $\int_{\ell_e}^{u_e} x \cdot f_e(x) dx < \infty$ .

**Definition 1.** A generalized semi-Markov process is a tuple  $(S, \mathcal{E}, \mathbf{E}, \text{Succ}, \alpha_0)$  where

- $S$  is a finite set of states,
- $\mathcal{E}$  is a finite set of events,
- $\mathbf{E} : S \rightarrow 2^\mathcal{E}$  assigns to each state  $s$  a set of events  $\mathbf{E}(s) \neq \emptyset$  scheduled to occur in  $s$ ,
- $\text{Succ} : S \times 2^\mathcal{E} \rightarrow \mathcal{D}(S)$  is the successor function, i.e. assigns a probability distribution specifying the successor state to each state and set of events that occur simultaneously in this state, and
- $\alpha_0 \in \mathcal{D}(S)$  is the initial distribution.

A *configuration* is a pair  $(s, \nu)$  where  $s \in S$  and  $\nu$  is a *valuation* which assigns to every event  $e \in \mathbf{E}(s)$  the amount of time that elapsed since the event  $e$  was scheduled.<sup>1</sup> For convenience, we define  $\nu(e) = \perp$  whenever  $e \notin \mathbf{E}(s)$ , and we denote by  $\nu(\Delta)$  the amount of time spent in the previous configuration (initially, we put  $\nu(\Delta) = 0$ ). When a set of events  $E$  occurs and the process moves from  $s$  to a state  $s'$ , the valuation of *old* events of  $\mathbf{E}(s) \setminus \mathbf{E}(s')$  is discarded to  $\perp$ , the valuation of each *inherited* event of  $(\mathbf{E}(s') \cap \mathbf{E}(s)) \setminus E$  is increased by the time spent in  $s$ , and the valuation of each *new* event of  $(\mathbf{E}(s') \setminus \mathbf{E}(s)) \cup (\mathbf{E}(s') \cap E)$  is set to 0.

We illustrate the dynamics of GSMP on the example of Figure 1. Let the bounds of the fixed-delay events *roundtrip<sub>d</sub>*, *polling<sub>d</sub>*, and *stable<sub>d</sub>* be

<sup>1</sup> Usually, the valuation is defined to store the time left before the event appears. However, our definition is equivalent and more convenient for the general setting where both bounded and unbounded events appear.

1, 90, and 100, respectively. We start in the state *Idle*, i.e. in the configuration  $(Idle, ((polling\_d, 0), (stable\_d, 0), (\Delta, 0)))$  denoting that  $v(polling\_d) = 0$ ,  $v(stable\_d) = 0$ ,  $v(\Delta) = 0$ , and  $\perp$  is assigned to all other events. After 90 time units, the event *polling\_d* occurs and we move to  $(Init, ((query, 0), (stable\_d, 90), (\Delta, 90)))$ . Assume that the event *query* occurs in the state *Init* after 0.6 time units and we move to  $(Q-sent, ((response, 0), (roundtrip\_d, 0), (stable\_d, 90.6), (\Delta, 0.6)))$  and so forth.

A formal semantics of GSMP is usually defined in terms of general state-space Markov chains (GSSMC, see, e.g., [23]). A GSSMC is a stochastic process  $\Phi$  over a measurable state-space  $(\Gamma, \mathcal{G})$  whose dynamics is determined by an initial measure  $\mu$  on  $(\Gamma, \mathcal{G})$  and a *transition kernel*  $P$  which specifies one-step transition probabilities.<sup>2</sup> A given GSMP induces a GSSMC whose state-space consists of all configurations, the initial measure  $\mu$  is induced by  $\alpha_0$  in a natural way, and the transition kernel is determined by the dynamics of GSMP described above. Formally,

- $\Gamma$  is the set of all configurations, and  $\mathcal{G}$  is a  $\sigma$ -field over  $\Gamma$  induced by the discrete topology over  $S$  and the Borel  $\sigma$ -field over the set of all valuations;
- the initial measure  $\mu$  allows to start in configurations with zero valuation only, i.e. for  $A \in \mathcal{G}$  we have  $\mu(A) = \sum_{s \in Zero(A)} \alpha_0(s)$  where  $Zero(A) = \{s \in S \mid (s, \mathbf{0}) \in A\}$ ;
- the transition kernel  $P(z, A)$  describing the probability to move in one step from a configuration  $z = (s, v)$  to any configuration in a set  $A$  is defined as follows. It suffices to consider  $A$  of the form  $\{s'\} \times X$  where  $X$  is a measurable set of valuations. Let  $V$  and  $F$  be the sets of variable-delay and fixed-delay events, respectively, that are scheduled in  $s$ . Let  $F' \subseteq F$  be the set of fixed-delay events that can occur as first among the fixed-delay event enabled in  $z$ , i.e. that have in  $v$  the minimal remaining time  $u$ . Note that two variable-delay events occur simultaneously with probability zero. Hence, we consider all combinations of  $e \in V$  and  $t \in \mathbb{R}_{\geq 0}$  stating that

$$P(z, A) = \begin{cases} \sum_{e \in V} \int_0^\infty \text{Hit}(\{e\}, t) \cdot \text{Win}(\{e\}, t) dt & \text{if } F = \emptyset \\ \sum_{e \in V} \int_0^u \text{Hit}(\{e\}, t) \cdot \text{Win}(\{e\}, t) dt + \text{Hit}(F', u) \cdot \text{Win}(F', u) & \text{otherwise,} \end{cases}$$

where the term  $\text{Hit}(E, t)$  denotes the conditional probability of hitting  $A$  under the condition that  $E$  occurs at time  $t$  and the term  $\text{Win}(E, t)$  denotes the probability (density) of  $E$  occurring at time  $t$ . Formally,

$$\text{Hit}(E, t) = \text{Succ}(s, E)(s') \cdot \mathbf{1}[v' \in X]$$

where  $\mathbf{1}[v' \in X]$  is the indicator function and  $v'$  is the valuation after the transition, i.e.  $v'(e)$  is  $\perp$ , or  $v(e) + t$ , or 0 for each old, or inherited, or new event  $e$ , respectively; and  $v'(\Delta) = t$ . The most complicated part is the definition of  $\text{Win}(E, t)$  which intuitively corresponds to the probability that  $E$  is the set of events “winning” the competition among the events scheduled in  $s$  at time  $t$ . First, we define a “shifted” density function  $f_{e|v(e)}$  that takes into account that the time  $v(e)$  has already elapsed. Formally, for a variable-delay event  $e$  and any elapsed time  $v(e) < u_e$ , we define

$$f_{e|v(e)}(x) = \frac{f_e(x + v(e))}{\int_{v(e)}^\infty f_e(y) dy} \quad \text{if } x \geq 0.$$

<sup>2</sup> Precisely, transition kernel is a function  $P : \Gamma \times \mathcal{G} \rightarrow [0, 1]$  such that  $P(z, \cdot)$  is a probability measure over  $(\Gamma, \mathcal{G})$  for each  $z \in \Gamma$ ; and  $P(\cdot, A)$  is a measurable function for each  $A \in \mathcal{G}$ .

Otherwise, we define  $f_{e|v(e)}(x) = 0$ . The denominator scales the function so that  $f_{e|v(e)}$  is again a density function. Finally,

$$\text{Win}(E, t) = \begin{cases} f_{e|v(e)}(t) \cdot \prod_{c \in V \setminus E} \int_t^\infty f_{c|v(c)}(y) dy & \text{if } E = \{e\} \subseteq V \\ \prod_{c \in V} \int_t^\infty f_{c|v(c)}(y) dy & \text{if } E = F' \subseteq F \\ 0 & \text{otherwise.} \end{cases}$$

A *run* of the Markov chain is an infinite sequence  $\sigma = z_0 z_1 z_2 \dots$  of configurations. The Markov chain is defined on the probability space  $(\Omega, \mathcal{F}, \mathcal{P})$  where  $\Omega$  is the set of all runs,  $\mathcal{F}$  is the product  $\sigma$ -field  $\bigotimes_{i=0}^\infty \mathcal{G}$ , and  $\mathcal{P}$  is the unique probability measure such that for every finite sequence  $A_0, \dots, A_n \in \mathcal{G}$  we have that

$$\mathcal{P}(\Phi_0 \in A_0, \dots, \Phi_n \in A_n) = \int_{z_0 \in A_0} \dots \int_{z_{n-1} \in A_{n-1}} \mu(dz_0) \cdot P(z_0, dz_1) \dots P(z_{n-1}, A_n)$$

where each  $\Phi_i$  is the  $i$ -th projection of an element in  $\Omega$  (the  $i$ -th configuration of a run).

Finally, we define an  $m$ -step transition kernel  $P^m$  inductively as  $P^1(z, A) = P(z, A)$  and  $P^{i+1}(z, A) = \int_{\Gamma} P(z, dy) \cdot P^i(y, A)$ .

## 2.2 Frequency measures

Our attention focuses on frequencies of a fixed state  $\mathring{s} \in S$  in the runs of the Markov chain. Let  $\sigma = (s_0, v_0) (s_1, v_1) \dots$  be a run. We define


$$\mathbf{d}(\sigma) = \lim_{n \rightarrow \infty} \frac{\sum_{i=0}^n \delta(s_i)}{n} \quad \mathbf{c}(\sigma) = \lim_{n \rightarrow \infty} \frac{\sum_{i=0}^n \delta(s_i) \cdot v_{i+1}(\Delta)}{\sum_{i=0}^n v_{i+1}(\Delta)}$$

where  $\delta(s_i)$  is equal to 1 when  $s_i = \mathring{s}$ , and 0 otherwise. We recall that  $v_{i+1}(\Delta)$  is the time spent in state  $s_i$  before moving to  $s_{i+1}$ . We say that the random variable  $\mathbf{d}$  or  $\mathbf{c}$  is *well-defined* for a run  $\sigma$  if the corresponding limit exists for  $\sigma$ . Then,  $\mathbf{d}$  corresponds to the frequency of discrete visits to the state  $\mathring{s}$  and  $\mathbf{c}$  corresponds to the ratio of time spent in the state  $\mathring{s}$ .

## 2.3 Region graph

In order to state the results in a simpler way, we introduce the *region graph*, a standard notion from the area of timed automata [5]. It is a finite partition of the uncountable set of configurations. First, we define the region relation  $\sim$ . For  $a, b \in \mathbb{R}$ , we say that  $a$  and  $b$  *agree on integral part* if  $\text{int}(a) = \text{int}(b)$  and neither or both  $a, b$  are integers. Further, we set the bound  $B = \max(\{\ell_e, u_e \mid e \in \mathcal{E}\} \setminus \{\infty\})$ . Finally, we put  $(s_1, v_1) \sim (s_2, v_2)$  if

- $s_1 = s_2$ ;
- for all  $e \in \mathbf{E}(s_1)$  we have that  $v_1(e)$  and  $v_2(e)$  agree on integral parts or are both greater than  $B$ ;
- for all  $e, f \in \mathbf{E}(s_1)$  with  $v_1(e) \leq B$  and  $v_1(f) \leq B$  we have that  $\text{frac}(v_1(e)) \leq \text{frac}(v_1(f))$  iff  $\text{frac}(v_2(e)) \leq \text{frac}(v_2(f))$ .



**Fig. 2.** A GSMP of a producer-consumer system. The events  $p$ ,  $t$ , and  $c$  model that a packet production, transport, and consumption is finished, respectively. Below each state label, there is the set of scheduled events. The fixed-delay events  $p$  and  $c$  have  $l_p = u_p = l_c = u_c = 1$  and the uniformly distributed variable-delay event  $t$  has  $l_t = 0$  and  $u_t = 1$ .

Note that  $\sim$  is an equivalence with finite index. The equivalence classes of  $\sim$  are called *regions*. We define a finite *region graph*  $G = (V, E)$  where the set of vertices  $V$  is the set of regions and for every pair of regions  $R, R'$  there is an edge  $(R, R') \in E$  iff  $P(z, R') > 0$  for some  $z \in R$ . The construction is correct because all states in the same region have the same one-step qualitative behavior (for details, see Appendix B.1).

### 3 Two fixed-delay events

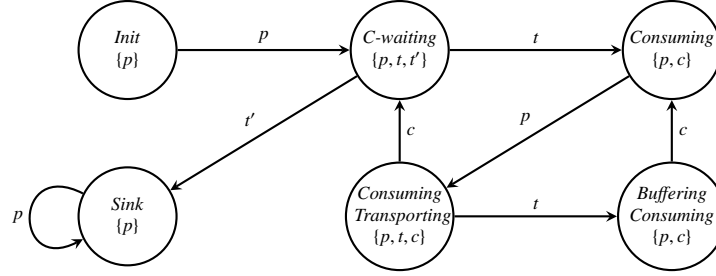
Now, we explain in more detail what problems can be caused by fixed-delay events. We start with an example of a GSMP with two fixed-delay events for which it is not true that the variables  $\mathbf{d}$  and  $\mathbf{c}$  are well-defined for almost all runs. Then we show some other examples of GSMP with fixed-delay events that disprove some results from literature. In the next section, we provide positive results when the number and type of fixed-delay events are limited.

#### When the frequencies $\mathbf{d}$ and $\mathbf{c}$ are not well-defined

In Figure 2, we show an example of a GSMP with two fixed-delay events and one variable-delay event for which it is not true that the variables  $\mathbf{d}$  and  $\mathbf{c}$  are well-defined for almost all runs. It models the following producer-consumer system. We use three components – a producer, a transporter and a consumer of packets. The components work in parallel but each component can process (i.e. produce, transport, or consume) at most one packet at a time.

Consider the following time requirements: each packet production takes *exactly* 1 time unit, each transport takes *at most* 1 time unit, and each consumption takes again *exactly* 1 time unit. As there are no limitations to block the producer, it is working for all the time and new packets are produced precisely each time unit. As the transport takes shorter time than the production, every new packet is immediately taken by the transporter and no buffer is needed at this place. When a packet arrives to the consumer, the consumption is started immediately if the consumer is waiting; otherwise, the packet





**Fig. 3.** A GSMP with two fixed-delay events  $p$  and  $c$  (with  $l_p = u_p = l_c = u_c = 1$ ), a uniformly distributed variable-delay events  $t, t'$  (with  $l_t = l_{t'} = 0$  and  $u_t = u_{t'} = 1$ ).

is stored into a buffer. When the consumption is finished and the buffer is empty, the consumer waits; otherwise, a new consumption starts immediately.

In the GSMP in Figure 2, the consumer has two modules – one is in operation and the other idles at a time – when the consumer enters the waiting state, it switches the modules. The labels 1 and 2 denote which module of the consumer is in operation.

One can easily observe that the consumer enters the waiting state (and switches the modules) if and only if the current transport takes more time than it has ever taken. As the transport time is bounded by 1, it gets harder and harder to break the record. As a result, the system stays in the current module on average for longer time than in the previous module. Therefore, due to the successively prolonging stays in the modules, the frequencies for 1-states and 2-states oscillate. For precise computations, see Appendix A.1. We conclude the above observation by the following theorem.

**Theorem 1.** *There is a GSMP (with two fixed-delay events and one variable-delay event) for which it is not true that the variables  $\mathbf{c}$  and  $\mathbf{d}$  are almost surely well-defined.*

### Counterexamples

In [3,4] there are algorithms for GSMP model checking based on the region construction. They rely on two crucial statements of the papers:

1. Almost all runs end in some of the bottom strongly connected components (BSCC) of the region graph.
2. Almost all runs entering a BSCC visit all regions of the component infinitely often.

Both of these statements are true for finite state Markov chains. In the following, we show that neither of them has to be valid for region graphs of GSMP.

Let us consider the GSMP depicted in Figure 3. This is a producer-consumer model similar to the previous example but we have only one module of the consumer here. Again, entering the state *C-waiting* indicates that the current transport takes more time than it has ever taken. In the state *C-waiting*, an additional event  $t'$  can occur and move the system into a state *Sink*. One can intuitively observe that we enter the state *C-waiting* less and less often and stay there for shorter and shorter time. Hence, the probability

that the event  $t'$  occurs in the state  $C$ -waiting is decreasing during the run. For precise computations proving the following claim, see Appendix A.2.

*Claim.* The probability to reach  $Sink$  from  $Init$  is strictly less than 1.

The above claim directly implies the following theorem thus disproving statement 1.

**Theorem 2.** *There is a GSMP (with two fixed-delay and two variable delay events) where the probability to reach any BSCC of the region graph is strictly smaller than 1.*

Now consider in Figure 3 a transition under the event  $p$  from the state  $Sink$  to the state  $Init$  instead of the self-loop. This turns the whole region graph into a single BSCC. We prove that the state  $Sink$  is almost surely visited only finitely often. Indeed, let  $p < 1$  be the original probability to reach  $Sink$  guaranteed by the claim above. The probability to reach  $Sink$  from  $Sink$  again is also  $p$  as the only transition leading from  $Sink$  enters the initial configuration. Therefore, the probability to reach  $Sink$  infinitely often is  $\lim_{n \rightarrow \infty} p^n = 0$ . This proves the following theorem. Hence, the statement 2 of [3,4] is disproved, as well.

**Theorem 3.** *There is a GSMP (with two fixed-delay and two variable delay events) with strongly connected region graph and with a region that is reached infinitely often with probability 0.*

## 4 Single-ticking GSMP

First of all, motivated by the previous counterexamples, we identify the behavior of the fixed-delay events that may cause  $\mathbf{d}$  and  $\mathbf{c}$  to be undefined. The problem lies in fixed-delay events that can immediately schedule themselves whenever they occur; such an event can occur periodically like ticking of clocks. In the example of Figure 3, there are two such events  $p$  and  $c$ . The phase difference of their ticking gets smaller and smaller, causing the unstable behavior.

For two fixed-delay events  $e$  and  $e'$ , we say that  $e$  causes  $e'$  if there are states  $s, s'$  and a set of events  $E$  such that  $\text{Succ}(s, E)(s') > 0$ ,  $e \in E$ , and  $e'$  is newly scheduled in  $s'$ .

**Definition 2.** *A GSMP is called single-ticking if either there is no fixed-delay event or there is a strict total order  $<$  on fixed-delay events with the least element  $e$  (called ticking event) such that whenever  $f$  causes  $g$  then either  $f < g$  or  $f = g = e$ .*

From now on we restrict to single-ticking GSMP and prove our main positive result.

**Theorem 4.** *In single-ticking GSMP, the random variables  $\mathbf{d}$  and  $\mathbf{c}$  are well-defined for almost every run and admit only finitely many values. Precisely, almost every run reaches a BSCC of the region graph and for each BSCC  $B$  there are values  $d, c \in [0, 1]$  such that  $\mathbf{d}(\sigma) = d$  and  $\mathbf{c}(\sigma) = c$  for almost all runs  $\sigma$  that reach the BSCC  $B$ .*

The rest of this section is devoted to the proof of Theorem 4. First, we show that almost all runs end up trapped in some BSCC of the region graph. Second, we solve the problem while restricting to runs that start in a BSCC (as the initial part of a run outside of any BSCC is not relevant for the long run average behavior). We show that in a BSCC, the variables  $\mathbf{d}$  and  $\mathbf{c}$  are almost surely constant. The second part of the proof relies on several standard results from the theory of general state space Markov chains. Formally, the proof follows from Propositions 1 and 2 stated below.

#### 4.1 Reaching a BSCC

**Proposition 1.** *In single-ticking GSMP, almost every run reaches a BSCC of the region graph.*

The proof uses similar methods as the proof in [4]. By definition, the process moves along the edges of the region graph. From every region, there is a minimal path through the region graph into a BSCC, let  $n$  be the maximal length of all such paths. Hence, in at most  $n$  steps the process reaches a BSCC with positive probability from any configuration. Observe that if this probability was bounded from below, we would eventually reach a BSCC from any configuration almost surely. However, this probability can be arbitrarily small. Consider the following example with event  $e$  uniform on  $[0, 1]$  and event  $f$  uniform on  $[2, 3]$ . In an intuitive notation, let  $R$  be the region  $[0 < e < f < 1]$ . What is the probability that the event  $e$  occurs after the elapsed time of  $f$  reaches 1 (i.e. that the region  $[e = 0; 1 < f < 2]$  is reached)? For a configuration in  $R$  with valuation  $((e, 0.2), (f, 0.7))$  the probability is 0.5 but for another configuration in  $R$  with  $((e, 0.2), (f, 0.21))$  it is only 0.01. Notice that the transition probabilities depend on the difference of the fractional values of the clocks, we call this difference *separation*. Observe that in other situations, the separation of clocks from value 0 also matters.

**Definition 3.** *Let  $\delta > 0$ . We say that a configuration  $(s, \nu)$  is  $\delta$ -separated if for every  $x, y \in \{0\} \cup \{\nu(e) \mid e \in \mathbf{E}(s)\}$ , we have either  $|\text{frac}(x) - \text{frac}(y)| > \delta$  or  $\text{frac}(x) = \text{frac}(y)$ .*

We fix a  $\delta > 0$ . To finish the proof using the concept of  $\delta$ -separation, we need two observations. First, from *any* configuration we reach in  $m$  steps a  $\delta$ -separated configuration with probability at least  $q > 0$ . Second, the probability to reach a fixed region from *any*  $\delta$ -separated configuration is bounded from below by some  $p > 0$ . By repeating the two observations ad infinitum, we reach some BSCC almost surely. Let us state the claims. For proofs, see Appendix B.2.

**Lemma 1.** *There is  $\delta > 0$ ,  $m \in \mathbb{N}$  and  $q > 0$  such that from every configuration we reach a  $\delta$ -separated configuration in  $m$  steps with probability at least  $q$ .*

**Lemma 2.** *For every  $\delta > 0$  and  $k \in \mathbb{N}$  there is  $p > 0$  such that for any pair of regions  $R, R'$  connected by a path of length  $k$  and for any  $\delta$ -separated  $z \in R$ , we have  $P^k(z, R') > p$ .*

Lemma 2 holds even for unrestricted GSMP. Notice that Lemma 1 does not. As in the example of Figure 3, the separation may be non-increasing for all runs.

#### 4.2 Frequency in a BSCC

From now on, we deal with the bottom strongly connected components that are reached almost surely. Hence, we assume that the region graph  $G$  is strongly connected. We have to allow an arbitrary initial configuration  $z_0 = (s, \nu)$ ; in particular,  $\nu$  does not have to be a zero vector.<sup>3</sup>

<sup>3</sup> Technically, the initial measure is  $\mu(A) = 1$  if  $z_0 \in A$  and  $\mu(A) = 0$ , otherwise.

**Proposition 2.** *In a single-ticking GSMP with strongly connected region graph, there are values  $d, c \in [0, 1]$  such that for any initial configuration  $z_0$  and for almost all runs  $\sigma$  starting from  $z_0$ , we have that  $\mathbf{d}$  and  $\mathbf{c}$  are well-defined and  $\mathbf{d}(\sigma) = d$  and  $\mathbf{c}(\sigma) = c$ .*

We assume that the region graph is aperiodic in the following sense. A *period*  $p$  of a graph  $G$  is the greatest common divisor of lengths of all cycles in  $G$ . The graph  $G$  is *aperiodic* if  $p = 1$ . Under this assumption<sup>4</sup>, the chain  $\Phi$  is in some sense stable. Namely, (i)  $\Phi$  has a unique invariant measure that is independent of the initial measure and (ii) the strong law of large numbers (SLLN) holds for  $\Phi$ .

First, we show that (i) and (ii) imply the proposition. Let us recall the notions. We say that a probability measure  $\pi$  on  $(\Gamma, \mathcal{G})$  is *invariant* if for all  $A \in \mathcal{G}$

$$\pi(A) = \int_{\Gamma} \pi(dx) P(x, A).$$

The SLLN states that if  $h : \Gamma \rightarrow \mathbb{R}$  satisfies  $E_{\pi}[h] < \infty$ , then almost surely

$$\lim_{n \rightarrow \infty} \frac{\sum_{i=1}^n h(\Phi_i)}{n} = E_{\pi}[h], \quad (1)$$

where  $E_{\pi}[h]$  is the expected value of  $h$  according to the invariant measure  $\pi$ .

We set  $h$  as follows. For a run  $(s_0, v_0)(s_1, v_1) \cdots$ , let  $h(\Phi_i) = 1$  if  $s_i = \mathring{s}$  and 0, otherwise. We have  $E_{\pi}[h] < \infty$  since  $h \leq 1$ . From (1) we obtain that almost surely

$$\mathbf{d} = \lim_{n \rightarrow \infty} \frac{\sum_{i=1}^n h(\Phi_i)}{n} = E_{\pi}[h].$$

As a result,  $\mathbf{d}$  is well-defined and equals the constant value  $E_{\pi}[h]$  for almost all runs. We treat the variable  $\mathbf{c}$  similarly. Let  $W((s, v))$  denote the expected waiting time of the GSMP in the configuration  $(s, v)$ . We use a function  $\tau((s, v)) = W((s, v))$  if  $s = \mathring{s}$  and 0, otherwise. Since all the events have finite expectation, we have  $E_{\pi}[W] < \infty$  and  $E_{\pi}[\tau] < \infty$ . Furthermore, we show in Appendix B.3 that almost surely

$$\mathbf{c} = \lim_{n \rightarrow \infty} \frac{\sum_{i=1}^n \tau(\Phi_i)}{\sum_{i=1}^n W(\Phi_i)} = \frac{E_{\pi}[\tau]}{E_{\pi}[W]}.$$

Therefore,  $\mathbf{c}$  is well-defined and equals the constant  $E_{\pi}[\tau]/E_{\pi}[W]$  for almost all runs.

Second, we prove (i) and (ii). A standard technique of general state space Markov chains (see, e.g., [23]) yields (i) and (ii) for chains that satisfy the following condition. Roughly speaking, we search for a set of configurations  $C$  that is visited infinitely often and for some  $\ell$  the measures  $P^{\ell}(x, \cdot)$  and  $P^{\ell}(y, \cdot)$  are very similar for any  $x, y \in C$ . This is formalized by the following lemma.

**Lemma 3.** *There is a measurable set of configurations  $C$  such that*

1. *there is  $k \in \mathbb{N}$  and  $\alpha > 0$  such that for every  $z \in \Gamma$  we have  $P^k(z, C) \geq \alpha$ , and*

---

<sup>4</sup> If the region graph has period  $p > 1$ , we can employ the standard technique and decompose the region graph (and the Markov chain) into  $p$  aperiodic components. The results for individual components yield straightforwardly the results for the whole Markov chain, see, e.g., [9].

2. there is  $\ell \in \mathbb{N}$ ,  $\beta > 0$ , and a probability measure  $\kappa$  such that for every  $z \in C$  and  $A \in \mathcal{G}$  we have  $P^\ell(z, A) \geq \beta \cdot \kappa(A)$ .

*Proof (Sketch).* Let  $e$  be the ticking event and  $R$  some reachable region where  $e$  is the event closest to its upper bound. We fix a sufficiently small  $\delta > 0$  and choose  $C$  to be the set of  $\delta$ -separated configurations of  $R$ . We prove the first part of the lemma similarly to Lemmata 1 and 2. As regards the second part, we define the measure  $\kappa$  uniformly on a hypercube  $X$  of configurations  $(s, \nu)$  that have  $\nu(e) = 0$  and  $\nu(f) \in (0, \delta)$ , for  $f \neq e$ . First, assume that  $e$  is the only fixed-delay event. We fix  $z = (s', \nu')$  in  $R$ ; let  $d = u_e - \nu'(e) > \delta$  be the time left in  $z$  before  $e$  occurs. For simplicity, we assume that each variable-delay events can occur after an arbitrary delay  $x \in (d - \delta, d)$ . Precisely, that it can occur in an  $\varepsilon$ -neighborhood of  $x$  with probability bounded from below by  $\beta \cdot \varepsilon$  where  $\beta$  is the minimal density value of all  $\mathcal{E}$ . Note that the variable-delay events can be “placed” this way arbitrarily in  $(0, \delta)$ . Therefore, when  $e$  occurs, it has value 0 and all variable-delay events can be in interval  $(0, \delta)$ . In other words, we have  $P^\ell(z, A) \geq \beta \cdot \kappa(A)$  for any measurable  $A \subseteq X$  and for  $\ell = |\mathcal{E}|$ .

Allowing other fixed-delay events causes some trouble because a fixed-delay event  $f \neq e$  cannot be “placed” arbitrarily. In the total order  $<$ , the event  $f$  can cause only strictly greater fixed-delay events. The greatest fixed-delay event can cause only variable-delay events that can be finally “placed” arbitrarily as described above.  $\square$

## 5 Approximations

In the previous section we have proved that in single-ticking GSMP,  $\mathbf{d}$  and  $\mathbf{c}$  are almost surely well-defined and for almost all runs they attain only finitely many values  $d_1, \dots, d_k$  and  $c_1, \dots, c_k$ , respectively. In this section we show how to approximate  $d_i$ ’s and  $c_i$ ’s and the probabilities that  $\mathbf{d}$  and  $\mathbf{c}$  attain these values, respectively.

**Theorem 5.** *In a single-ticking GSMP, let  $d_1, \dots, d_k$  and  $c_1, \dots, c_k$  be the discrete and timed frequencies, respectively, corresponding to BSCCs of the region graph. For all  $1 \leq i \leq k$ , the numbers  $d_i$  and  $c_i$  as well as the probabilities  $\mathcal{P}(\mathbf{d} = d_i)$  and  $\mathcal{P}(\mathbf{c} = c_i)$  can be approximated up to any  $\varepsilon > 0$ .*

*Proof.* Let  $X_1, \dots, X_k$  denote the sets of configurations in individual BSCCs and  $d_i$  and  $c_i$  correspond to  $X_i$ . Since we reach a BSCC almost surely, we have

$$\mathcal{P}(\mathbf{d} = d_i) = \sum_{j=1}^k \mathcal{P}(\mathbf{d} = d_i \mid \text{Reach}(X_j)) \cdot \mathcal{P}(\text{Reach}(X_j)) = \sum_{j=1}^k \mathbf{1}[d_j = d_i] \cdot \mathcal{P}(\text{Reach}(X_j))$$

where the second equality follows from the fact that almost all runs in the  $j$ -th BSCC yield the discrete frequency  $d_j$ . Therefore,  $\mathcal{P}(\mathbf{d} = d_i)$  and  $d_i$  can be approximated as follows using the methods of [24].

*Claim.* Let  $X$  be a set of all configurations in a BSCC  $\mathcal{B}$ ,  $X_{\hat{s}} \subseteq X$  the set of configurations with state  $\hat{s}$ , and  $d$  the frequency corresponding to  $\mathcal{B}$ . There are computable constants  $n_1, n_2 \in \mathbb{N}$  and  $p_1, p_2 > 0$  such that for every  $i \in \mathbb{N}$  and  $z_X \in X$  we have

$$\begin{aligned} |\mathcal{P}(\text{Reach}(X)) - P^i(z_0, X)| &\leq (1 - p_1)^{\lfloor i/n_1 \rfloor} \\ |d - P^i(z_X, X_{\hat{s}})| &\leq (1 - p_2)^{\lfloor i/n_2 \rfloor} \end{aligned}$$

Further, we want to approximate  $c_i = E_\pi[\tau]/E_\pi[W]$ , where  $\pi$  is the invariant measure on  $X_i$ . In other words, we need to approximate  $\int_{X_i} \tau(x)\pi(dx)$  and  $\int_{X_i} W(x)\pi(dx)$ . An  $n$ -th approximation  $w_n$  of  $E_\pi[W]$  can be gained by discretizing the part of the state space  $\{(s, \nu) \in \Gamma \mid \forall e \in \mathbf{E}(s) : \nu(e) \leq n\}$  into, e.g.,  $1/n$ -large hypercubes, where the invariant measure  $\pi$  is approximated using  $P^n$ . This approximation converges to  $E_\pi[W]$  since  $W$  is continuous and  $E_\pi[W]$  is finite. For the details of the following claim, see Appendix C.

*Claim.* On each region,  $W$  is continuous, and  $E_\pi[W]$  is finite.

This concludes the proof as  $\tau$  only differs from  $W$  in being identically zero on some regions; thus,  $E_\pi[\tau]$  can be approximated analogously.

## 6 Conclusions, future work

We have studied long run average properties of generalized semi-Markov processes with both fixed-delay and variable-delay events. We have shown that two or more (unrestricted) fixed-delay events lead to considerable complications regarding stability of GSMP. In particular, we have shown that the frequency of states of a GSMP may not be well-defined and that bottom strongly connected components of the region graph may not be reachable with probability one. This leads to counterexamples disproving several results from literature. On the other hand, for single-ticking GSMP we have proved that the frequencies of states are well-defined for almost all runs. Moreover, we have shown that almost every run has one of finitely many possible frequencies that can be effectively approximated (together with their probabilities) up to a given error tolerance.

In addition, the frequency measures can be easily extended into the mean payoff setting. Consider assigning real rewards to states. The mean payoff then corresponds to the frequency weighted by the rewards.

Concerning future work, the main issue is efficiency of algorithms for computing performance measures for GSMP. We plan to work on both better analytical methods as well as practicable approaches to Monte Carlo simulation. One may also consider extensions of our positive results to controlled GSMP and games on GSMP.

## References

1. de Alfaro, L.: How to specify and verify the long-run average behavior of probabilistic systems. In: Proceedings of LICS'98. pp. 454–465. IEEE Computer Society Press (1998)
2. Alur, R., Bernadsky, M.: Bounded model checking for GSMP models of stochastic real-time systems. In: Proceedings of 9th International Workshop Hybrid Systems: Computation and Control (HSCC). Lecture Notes in Computer Science, vol. 3927, pp. 19–33. Springer (2006)
3. Alur, R., Courcoubetis, C., Dill, D.: Model-checking for probabilistic real-time systems. In: Proceedings of ICALP'91. Lecture Notes in Computer Science, vol. 510, pp. 115–136. Springer (1991)
4. Alur, R., Courcoubetis, C., Dill, D.: Verifying automata specifications of probabilistic real-time systems. In: Real-Time: Theory in Practice. Lecture Notes in Computer Science, vol. 600, pp. 28–44. Springer (1992)
5. Alur, R., Dill, D.: A theory of timed automata. Theoretical Computer Science 126(2), 183–235 (1994)

6. Baier, C., Haverkort, B., Hermanns, H., Katoen, J.P.: Model-checking algorithms for continuous-time Markov chains. *IEEE Transactions on Software Engineering* 29(6), 524–541 (2003)
7. Barbot, B., Chen, T., Han, T., Katoen, J., Mereacre, A.: Efficient CTMC model checking of linear real-time objectives. *Tools and Algorithms for the Construction and Analysis of Systems* pp. 128–142 (2011)
8. Bernadsky, M., Alur, R.: Symbolic analysis for GSMP models with one stateful clock. In: *Proceedings of 10th International Workshop Hybrid Systems: Computation and Control (HSCC)*. Lecture Notes in Computer Science, vol. 4416, pp. 90–103. Springer (2007)
9. Brázdil, T., Krčál, J., Křetínský, J., Kučera, A., Řehák, V.: Measuring performance of continuous-time stochastic processes using timed automata. In: *Proceedings of 14th International Conference on Hybrid Systems: Computation and Control (HSCC'11)*. pp. 33–42. ACM Press (2011)
10. Ciardo, G., Jones III, R., Miner, A., Siminiceanu, R.: Logic and stochastic modeling with SMART. *Performance Evaluation* 63(6), 578–608 (2006)
11. D'Argenio, P., Katoen, J.: A theory of stochastic systems Part I: Stochastic automata. *Information and computation* 203(1), 1–38 (2005)
12. German, R., Lindemann, C.: Analysis of stochastic Petri nets by the method of supplementary variables. *Performance Evaluation* 20(1-3), 317–335 (1994)
13. Glynn, P.: A GSMP formalism for discrete event systems. *Proceedings of the IEEE* 77, 14–23 (1989)
14. Haas, P.: On simulation output analysis for generalized semi-markov processes. *Commun. Statist. Stochastic Models* 15, 53–80 (1999)
15. Haas, P.: *Stochastic Petri Nets: Modelling, Stability, Simulation*. Springer Series in Operations Research and Financial Engineering, Springer (2010)
16. Haas, P., Shedler, G.: Regenerative generalized semi-Markov processes. *Stochastic Models* 3(3), 409–438 (1987)
17. Lindemann, C., Reuys, A., Thummler, A.: The DSPNexpress 2.000 performance and dependability modeling environment. In: *Fault-Tolerant Computing, 1999. Digest of Papers. Twenty-Ninth Annual International Symposium on*. pp. 228–231. IEEE Computer Society Press (1999)
18. Lindemann, C., Shedler, G.: Numerical analysis of deterministic and stochastic Petri nets with concurrent deterministic transitions. *Performance Evaluation* 27, 565–582 (1996)
19. López, G., Hermanns, H., Katoen, J.: Beyond memoryless distributions: Model checking semi-Markov chains. *Process Algebra and Probabilistic Methods. Performance Modelling and Verification* pp. 57–70 (2001)
20. Marsan, M., Chiola, G.: On Petri nets with deterministic and exponentially distributed firing times. *Advances in Petri Nets 1987* pp. 132–145 (1987)
21. Marsan, M., Balbo, G., Conte, G., Donatelli, S., Franceschinis, G.: *Modelling with Generalized Stochastic Petri Nets*. Wiley (1995)
22. Matthes, K.: Zur Theorie der Bedienungsprozesse. *Transactions of the Third Prague Conference on Information Theory, Statistical Decision Functions, Random Processes* pp. 513–528 (1962)
23. Meyn, S., Tweedie, R.: *Markov Chains and Stochastic Stability*. Cambridge University Press (2009)
24. Roberts, G., Rosenthal, J.: General state space Markov chains and MCMC algorithms. *Probability Surveys* 1, 20–71 (2004)
25. Rosenthal, J.: *A first look at rigorous probability theory*. World Scientific Publishing (2006)
26. Younes, H., Simmons, R.: Probabilistic verification of discrete event systems using acceptance sampling. In: *Computer Aided Verification*. pp. 23–39. Springer (2002)

## A Details on counterexamples

**Definition 4.** A distance of two events  $e$  and  $f$  (in this order) in a configuration  $(s, \nu)$  is  $\text{frac}(\nu(f) - \nu(e))$ .

### A.1 When the frequencies $\mathbf{d}$ and $\mathbf{c}$ are not well-defined: Proof of Theorem 1

In the following, we prove that in our example  $\mathbf{d}$  and  $\mathbf{c}$  are not well-defined for almost all runs. Namely, that there is a set of runs with positive measure such that for these runs the partial sums oscillate.

After setting the initial distance of events  $p$  and  $c$ , every run stays in the 1-states (labeled with 1) until the distance is lessened in the state 2 *C-waiting*. This sojourn in the 1-states is called the first *phase*. Then the run continues with the second phase now in the 2-states until the distance is lessened again and it moves back to 1-states and begins the third phase etc. Each phase consists of repeating several *attempts*, i.e. running through the cycle of length three. In each attempt the distance gets smaller with probability  $d$  (where  $d$  is the current distance) and stays the same with probability  $1 - d$  due to the uniform distribution of  $t$ . This behavior corresponds to the geometric distribution. The density on the new distance is uniform on the whole  $d$ . A phase is called *strong* if the newly generated distance is at most half of the old one. Further, we define a *half-life* to be a maximum continuous sequence of phases where exactly the last one is strong. Every run can thus be uniquely decomposed into a sequence of half-lives. The random variable stating the distance at the beginning of the  $j$ -th phase of the  $i$ -th half-life is denoted  $D_{i,j}$ . Denoting the number of phases in the  $i$ -th half-life by  $L(i)$  we get  $D_{n-1,L(i)} \geq 2D_{n,1}$ . Thus by induction, we have for all  $n, i \in \mathbb{N}$  and  $j \leq L(n - i)$ ,

$$D_{n-i,j} \geq 2^i \cdot D_{n,1} \quad (2)$$

Further, let  $S_{i,j}$  be the number of attempts in the  $j$ -th phase of the  $i$ -th half-life, i.e. a *length* of this phase. We can now prove the following lemma. Roughly speaking, there are runs (of overall positive measure) where some phase is longer than the overall length of all phases up to that point. Note that the precise statement of the lemma implies moreover that this happens even infinitely often on runs of overall positive measure.

**Lemma 4.** *There are  $\alpha > 0$  and  $m > 0$ , such that for every  $n > 1$  there is a set  $\mathcal{R}_n$  of measure at least  $m$  of runs satisfying*

$$S_{n,1} \geq \alpha \sum_{\substack{i=1..n-1 \\ j=1..L(i)}} S_{i,j}$$

*Proof.* We set  $\alpha = 2/(3 \cdot (6 + 2 \cdot 3)) = 1/18$  and  $m = 1/4$  and let  $n > 1$  be arbitrary. We define the set  $\mathcal{R}_n$  to be the set of all runs  $\sigma$  such that the following conditions hold:

1.  $S_{n,1} > 1/(2D_{n,1})$ ,  
(the length of the “last” phase is above its expectation),
2. for all  $1 \leq i < n$ ,  $L(i) \leq (n - i) + 3$ ,  
(previous half-lives have no more phases than  $n + 2, n + 1, \dots, 5, 4$ , respectively),



3. for all  $1 \leq i < n$  and  $1 \leq j \leq L(i)$ ,  $S_{i,j} \leq 3(n-i)/D_{i,j}$ ,  
(all phases in previous half-lives are short w.r.t their expectations).

Denote  $D := D_{n,1}$ . We firstly prove that  $S_{n,1} \geq \alpha \sum_{i=1..n-1, j=1..L(i)} S_{i,j}$  for all runs in  $\mathcal{R}_n$ . Due to the inequality (2) and requirements 2. and 3., we can bound the overall length of all previous phases by

$$\sum_{\substack{i=1..n-1 \\ j=1..L(i)}} S_{i,j} \leq \sum_{i=1}^{n-1} \frac{(i+3) \cdot 3i}{2^i \cdot D} \leq \sum_{i=1}^{\infty} \frac{(i+3) \cdot 3i}{2^i \cdot D} = \frac{3(6+2 \cdot 3)}{D} = \frac{1}{2\alpha D}$$

and conclude by the requirement 1.

It remains to prove that measure of  $\mathcal{R}_n$  is at least  $m$ . We investigate the measures of the runs described by requirements 1.–3. Firstly, the probability that  $S_{n,1} > \frac{1}{2D_{n,1}}$  is  $(1 - D_{n,1})^{1/2D_{n,1}}$ , which approaches  $1/\sqrt{e}$  as  $n$  approaches infinity and is thus greater than  $1/2$  for  $D_{n,1} \leq 1/2$ , i.e. for  $n \geq 2$ . Out of this set of runs of measure  $1/2$  we need to cut off all runs that do not satisfy requirements 2. or 3. As for 2., the probability of  $i$ -th half-life failing to satisfy 2. is  $(1/2)^{(n-i)+3}$  corresponding to at least  $(n-i)+3$  successive non-strong phases. Therefore, 2. cuts off  $\sum_{i=1}^{n-1} 1/2^{(n-i)+3} = \sum_{i=1}^{n-1} 1/2^{i+3} \leq \sum_{i=1}^{\infty} 1/2^{i+3} = 1/2^3$ . From the remaining runs we need to cut off all runs violating 3. Since the probability of each  $S_{i,j}$  failing is  $(1 - D_{i,j})^{3(n-i)/D_{i,j}}$ , the overall probability of all violating runs is due to 2 at most

$$\begin{aligned} \sum_{i=1}^{n-1} \sum_{j=1}^{L(i)} (1 - D_{i,j})^{3(n-i)/D_{i,j}} &= \sum_{i=1}^{n-1} \sum_{j=1}^{L(n-i)} (1 - D_{n-i,j})^{3i/D_{n-i,j}} \leq \sum_{i=1}^{n-1} \sum_{j=1}^{L(n-i)} (1 - 2^i D)^{3i/2^i D} \\ &\leq \sum_{i=1}^{n-1} (i+3)(1 - 2^i D)^{3i/2^i D} \leq \sum_{i=1}^{\infty} (i+3)(1/e)^{3i} \\ &= \frac{4e^3 - 3}{(e^3 - 1)^2} < 1/4 \end{aligned}$$

Altogether the measure of  $\mathcal{R}_n$  is at least  $m = 1/2 - 1/8 - 1/4 = 1/8$ .  $\square$

Due to the previous lemma, moreover, there is a set  $\mathcal{R}$  of runs of positive measure such that each run of  $\mathcal{R}$  is contained in infinitely many  $\mathcal{R}_n$ 's.

Let us measure the frequency of 1-states (we slightly abuse the notation and denote by  $\mathbf{d}(\sigma)$  and  $\mathbf{c}(\sigma)$  the sum of frequencies of all 1-states instead of one single state  $\mathfrak{s}$ ). We prove that neither  $\mathbf{d}(\sigma)$  nor  $\mathbf{c}(\sigma)$  is well-defined on any  $\sigma \in \mathcal{R}$ . Since attempts last for one time unit, non-existence of  $\mathbf{d}(\sigma)$  implies non-existence of  $\mathbf{c}(\sigma)$ . Thus, assume for a contradiction that  $\mathbf{d}(\sigma)$  is well-defined. Denote  $s_i$  the number of attempts in the  $i$ -th phase. Because 1-states are visited exactly in odd phases, we have

$$\mathbf{d}(\sigma) = \lim_{n \rightarrow \infty} \frac{\sum_{i=1}^n s_i \cdot \text{odd}(i)}{\sum_{i=1}^n s_i}$$

where  $\text{odd}(i) = 1$  if  $i$  is odd and 0 otherwise. By the definition of limit, for every  $\varepsilon > 0$  there is  $n_0$  such that for all  $n > n_0$

$$\left| \frac{\sum_{i=1}^n s_i \cdot \text{odd}(i)}{\sum_{i=1}^n s_i} - \frac{\sum_{i=1}^{n-1} s_i \cdot \text{odd}(i)}{\sum_{i=1}^{n-1} s_i} \right| < \varepsilon \quad (3)$$

Due to the lemma,  $s_n \geq \alpha \sum_{i=1..n-1} s_i$  happens for infinitely many both odd and even phases  $n$  on  $\sigma \in \mathcal{R}$ . Now let  $\mathbf{d}(\sigma) \leq 1/2$ , the other case is handled symmetrically. Let  $\varepsilon$  be such that  $\alpha \geq \frac{\varepsilon}{1-2\varepsilon-\mathbf{d}(\sigma)}$ , and we choose an odd  $n > n_0$  satisfying  $s_n \geq \alpha \sum_{i=1}^{n-1} s_i \geq \frac{\varepsilon}{1-2\varepsilon-\mathbf{d}(\sigma)} \sum_{i=1}^{n-1} s_i$ . Denoting  $A = \sum_{i=1}^{n-1} s_i$  and  $O = \sum_{i=1}^{n-1} s_i \cdot \text{odd}(i)$  we get from (3) that

$$\begin{aligned} \frac{O + s_n}{A + s_n} - \frac{O}{A} &\geq \frac{O + \frac{\varepsilon}{1-2\varepsilon-\mathbf{d}(\sigma)} A}{A + \frac{\varepsilon}{1-2\varepsilon-\mathbf{d}(\sigma)} A} - \frac{O}{A} \stackrel{(*)}{=} \frac{\varepsilon}{1 - \mathbf{d}(\sigma) - \varepsilon} \cdot \left(1 - \frac{O}{A}\right) \\ &\stackrel{(**)}{\geq} \frac{\varepsilon}{1 - \mathbf{d}(\sigma) - \varepsilon} \cdot (1 - \mathbf{d}(\sigma) - \varepsilon) = \varepsilon \end{aligned}$$

which is a contradiction with (3). Notice that we omitted the absolute value from (3) because for an odd  $n$  the term is non-negative. The equality (\*) is a straightforward manipulation. In (\*\*) we use, similarly to (3), that  $|\frac{O}{A} - \mathbf{d}(\sigma)| < \varepsilon$ .

## A.2 Counterexamples: Proof of Claim

In the following, we prove that the probability to reach the state *Sink* is strictly less than 1.

Similarly as in the proof of Theorem 1, we introduce phases and half-lives and proceed with similar but somewhat simpler arguments. Let  $d$  be the distance of events  $p$  and  $c$ . Note that  $1-d$  is the maximum length of transportation so far. The initial distance is generated in the state *C-waiting* with a uniform distribution on  $(0, 1)$ . After that, the distance gets smaller and smaller over the time (if we ignore the states where the distance is not defined) whenever we enter the state *C-waiting*. Each sequence between two successive visits of *C-waiting* on a run is called a *phase* of this run. After each phase the current distance is lessened. The density on the new distance is uniform on the whole  $d$ . A phase is called *strong* if the newly generated distance is at most half of the old one. Further, we define a *half-life* to be a maximum continuous sequence of phases where exactly the last one is strong. Every run can thus be uniquely decomposed into a sequence of half-lives (with the last segment being possibly infinite if *C-waiting* is never reached again). The random variable stating the distance at the beginning of the  $i$ -th half-life is denoted by  $D_i$ . By definition,  $D_i \leq D_{i-1}/2$  and by induction, for every run with at least  $i$  half-lives

$$D_i \leq 1/2^i. \quad (4)$$

Denoting the number of phases in the  $i$ -th half-life by  $L(i)$ , we can prove the following lemma.

**Lemma 5.** *There is  $m > 0$  such that for every  $n > 1$  the set  $\mathcal{R}_n$  of runs  $\sigma$  satisfying*

1.  *$\sigma$  does not visit Sink during the first  $n$  half-lives, and*
2. *for every  $1 \leq i \leq n$  not exceeding the number of half-lives of  $\sigma$ ,  $L(i)(\sigma) \leq 2 \cdot i$*

*has measure at least  $m$ .*

This lemma concludes the proof, as there is a set of runs of measure at least  $m$  that never reach the state *Sink*. We now prove the lemma.

Firstly, for every  $n$  we bound the measure of runs satisfying the second condition. The probability that  $2i$  consecutive phases are not strong, i.e.  $L(i) > 2i$ , is  $1/2^{2i}$  as  $t$  is distributed uniformly. Therefore, the probability that there is  $i \leq n$  with  $L(i) > 2i$  is less than  $\sum_{i=1}^n 1/2^{2i}$ . This probability is thus for all  $n \in \mathbb{N}$  less than  $\sum_{i=1}^{\infty} 1/2^{2i} = 1/3$ . Hence, for each  $n$  at least  $2/3$  of runs satisfy the second condition.

Secondly, we prove that at least  $m'$  of runs satisfying the second condition also satisfy the first condition. This concludes the proof of the lemma as  $m'$  is independent of  $n$  (a precise computation reveals that  $m' > 0.009$ ).

Recall that  $D_i \leq 1/2^i$  and we assume that  $L(i) \leq 2i$ . Therefore, the probability that *Sink* is not reached during the  $i$ -th half-life is at least  $(1 - 1/2^i)^{2i}$  as  $t'$  is distributed uniformly and the distance can only get smaller during the half-life. Hence, the probability that in none of the first  $n$  half-lives *Sink* is reached is at least

$$\prod_{i=1}^n (1 - 1/2^i)^{2i}$$

Thus, for every  $n$ , the probability is greater than  $\prod_{i=1}^{\infty} (1 - 1/2^i)^{2i} =: m'$ . It remains to show that  $m' > 0$ . This is equivalent to  $\sum_{i=1}^{\infty} \ln(1 - 1/2^i)^{2i} > -\infty$ , which in turn can be rewritten as

$$2 \sum_{i=1}^{\infty} i \ln \left( \frac{2^i}{2^i - 1} \right) < \infty$$

Since  $\sum_{i=1}^{\infty} 1/i^2$  converges, it is sufficient to prove that

$$\ln \left( \frac{2^i}{2^i - 1} \right) \in O(1/i^3).$$

We get the result by rewriting the term in the form of an approximation of the derivative of  $\ln$  in  $2^i - 1$  which is smaller than the derivative of  $\ln$  in  $2^i - 1$  because  $\ln$  is concave

$$\ln \left( \frac{2^i}{2^i - 1} \right) = \frac{\ln(2^i) - \ln(2^i - 1)}{1} \leq \ln'(2^i - 1) = \frac{1}{2^i - 1} \in O(1/i^3).$$

□

## B Proofs of Section 4

In this section, by saying *value* of an event  $e$ , we mean the fractional part  $\text{frac}(v(e))$  when the valuation  $v$  is clear from context. Furthermore, by  $M$  we denote the sum of  $u_e$  of all fixed-delay events.

### B.1 Correctness of the region graph construction

The correctness of the region graph construction is based on the fact that configurations in one region can qualitatively reach the same regions in one step.

**Lemma 6.** *Let  $z \sim z'$  be configurations and  $R$  be a region. We have  $P(z, R) > 0$  iff  $P(z', R) > 0$ .*

*Proof.* For the sake of contradiction, let us fix a region  $R$  and a pair of configurations  $z \sim z'$  such that  $P(z, R) > 0$  and  $P(z', R) = 0$ . Let  $z = (s, \nu)$  and  $z' = (s, \nu')$ .

First, let us deal with the fixed-delay events. Let us assume that the part of  $P(z, R)$  contributed by the variable-delay events  $V$  is zero, i.e.  $\sum_{e \in V} \int_0^\infty \text{Hit}(\{e\}, t) \cdot \text{Win}(\{e\}, t) dt = 0$ . Then the set  $E$  of fixed-delay events scheduled with the minimal remaining time in  $z$  must be non-empty, i.e. some  $e \in E$ . We have

$$\begin{aligned} P(z, R) &= \text{Succ}(s, E)(s') \cdot \mathbf{1}[\bar{\nu} \in R] \cdot \prod_{c \in V} \int_{\nu(e)}^\infty f_{c|\nu(c)}(y) dy > 0 \\ P(z', R) &= \text{Succ}(s, E)(s') \cdot \mathbf{1}[\bar{\nu}' \in R] \cdot \prod_{c \in V} \int_{\nu'(e)}^\infty f_{c|\nu'(c)}(y) dy = 0 \end{aligned}$$

where  $s'$  is the control state of the region  $R$  and  $\bar{\nu}$  and  $\bar{\nu}'$  are the valuations after the transitions from  $z$  and  $z'$ , respectively. It is easy to see that from  $z \sim z'$  we get that  $\bar{\nu} \in R$  iff  $\bar{\nu}' \in R$ . Hence,  $P(z, R)$  and  $P(z', R)$  can only differ in the big product. Let us fix any  $c \in V$ . We show that  $\int_{\nu(e)}^\infty f_{c|\nu(c)}(y) dy$  is positive. Recall that the density function  $f_c$  can qualitatively change only on integral values. Both  $z$  and  $z'$  have the same order of events' values. Hence, the integral is positive for  $\nu'$  iff it is positive for  $\nu$ . We get  $P(z', R) > 0$  which is a contradiction.

On the other hand, let us assume that there is a variable-delay event  $e \in V$  such that

$$\int_0^\infty \text{Succ}(s, \{e\})(s') \cdot \mathbf{1}[\nu_t \in R] \cdot f_{e|\nu(e)}(t) \cdot \prod_{c \in V \setminus \{e\}} \int_t^\infty f_{c|\nu(c)}(y) dy dt > 0$$

where  $\nu_t$  is the valuation after the transition from  $z$  with waiting time  $t$ . There must be an interval  $I$  such that for every  $t \in I$  we have that  $f_{e|\nu(e)}(t)$  is positive,  $\mathbf{1}[\nu_t \in R] = 1$ , and  $\int_t^\infty f_{c|\nu(c)}(y) dy > 0$  for any  $c \in V \setminus \{e\}$ . From the definition of the region relation, this interval  $I$  corresponds to an interval between two adjacent events in  $\nu$ . Since  $z \sim z'$ , there must be also an interval  $I'$  such that for every  $t \in I'$  we have that  $f_{e|\nu'(e)}(t)$  is positive,  $\mathbf{1}[\nu'_t \in R] = 1$ , and  $\int_t^\infty f_{c|\nu'(c)}(y) dy > 0$  for any  $c \in V \setminus \{e\}$ . Hence,  $P(z', R) > 0$ , contradiction.  $\square$

## B.2 Proof of Proposition 1

**Lemma 1.** There is  $\delta > 0$ ,  $m \in \mathbb{N}$  and  $q > 0$  such that from every configuration we reach a  $\delta$ -separated configuration in  $m$  steps with probability at least  $q$ .

*Proof.* We divide the  $[0, 1]$  line segment into  $3 \cdot |\mathcal{E}| + 1$  slots of equal length  $\delta$ . Each value of a scheduled event lies in some slot. We show how to reach a configuration where the values are separated by empty slots.

As the time flows, the values shift along the slots. When an event occurs, values of all the newly scheduled events are placed to 0. The variable-delay events can be easily separated if we guarantee that variable-delay events occur in an interval of time when the first and the last slots of the line segment are empty.

We let the already scheduled variable-delay events occur arbitrarily. For each newly scheduled variable-delay event we place a token at the end of an empty slot with its left and right neighbour slots empty as well (i.e. there is no clock's value nor any other token in these three slots). Such slot must always exist since there are more slots than  $3 \cdot |\mathcal{E}|$ . As the time flows we move the tokens along with the events' values. Whenever a token reaches 1 on the  $[0, 1]$  line segment, we do the following. If the valuation of its associated event is not between its lower and upper bound, we move the token to 0 and wait one more time unit. Otherwise, we let the associated event occur from now up to time  $\delta$ . Indeed, for any moment in this interval, the first and the last slots of the line segment are empty. The probability that all variable-delay events occur in these prescribed intervals is bounded from below because events' densities are bounded from below.

The fixed-delay events cause more trouble because they occur at a fixed moment; possibly in an occupied slot. If a fixed-delay event always schedules itself (or there is a cycle of fixed-delay events that schedule each other), its value can never be separated from another such fixed-delay event. Therefore, we have limited ourselves to at most one ticking event  $e$ . Observe that every other event has its lifetime – the length of the chain of fixed-delay events that schedule each other. The lifetime of any fixed-delay event is obviously bounded by  $M$  which is the sum of delay of all fixed-delay events in the system. After time  $M$ , all the old non-ticking events “die”, all the newly scheduled non-ticking events are separated because they are initially scheduled by a variable-delay event. Therefore, we let the variable-delay events occur as explained above for  $m$  steps such that it takes more than  $M$  time units in total. We set  $m = \lceil M/\delta \rceil$  since each step takes at least  $\delta$  time.  $\square$

**Lemma 2.** For every  $\delta > 0$  and  $k \in \mathbb{N}$  there is  $p > 0$  such that for any pair of regions  $R, R'$  connected by a path of length  $k$  and for any  $\delta$ -separated  $z \in R$ , we have  $P^k(z, R') > p$ . Furthermore,  $P^k(z, X) > p$  where  $X \subseteq R'$  is the set of  $(\delta/3^k)$ -separated configurations.

*Proof.* Let  $z \in R_0$ ,  $k \in \mathbb{N}$ , and  $R_0, R_1, \dots, R_k$  be a path in the region graph to the region  $R = R_k$ . We can follow this path so that in each step we lose two thirds of the separation. At last, we reach a  $(\delta/3^k)$ -separated configuration in the target region  $R_k$ . We get the overall bound on probabilities from bounds on every step.

In each step either a variable-delay event or a set of fixed-delay events occur. Let  $\delta'$  be the separation in the current step. To follow the region path, a specified event must occur in an interval between two specified values which are  $\delta'$ -separated. A fixed-delay event occurs in this interval for sure because it has been scheduled this way. For a variable-delay event, we divide this interval into thirds and let the event occur in the middle subinterval. This happens with a probability bounded from below because events' densities are bounded from below. Furthermore, to follow the path in the region graph, no other event can occur sooner. Every other event has at least  $\delta'/3$  to its upper bound; the probability that it does *not* occur is again bounded from below.  $\square$

### B.3 Proof of Proposition 2

**Proposition 2.** In a single-ticking GSMP with strongly connected region graph, there are values  $d, c \in [0, 1]$  such that for any initial configuration  $z_0$  and for almost all runs  $\sigma$  starting from  $z_0$ , we have that  $\mathbf{d}$  and  $\mathbf{c}$  are well-defined and  $\mathbf{d}(\sigma) = d$  and  $\mathbf{c}(\sigma) = c$ .

*Proof.* First, we show using the following lemma that  $\Phi$  has a unique invariant measure and that the Strong Law of Large Numbers holds for  $\Phi$ . We prove the lemma later in this subsection.

**Lemma 3.** There is a measurable set of configurations  $C$  such that

1. there is  $k \in \mathbb{N}$  and  $\alpha > 0$  such that for every  $z \in \Gamma$  we have  $P^k(z, C) \geq \alpha$ , and
2. there is  $\ell \in \mathbb{N}$ ,  $\beta > 0$ , and a probability measure  $\kappa$  such that for every  $z \in C$  and  $A \in \mathcal{G}$  we have  $P^\ell(z, A) \geq \beta \cdot \kappa(A)$ .

A direct corollary of Lemma 3 is that the set of configurations is *small*.

**Definition 5.** Let  $n \in \mathbb{N}$ ,  $\varepsilon > 0$ , and  $\kappa$  be a probability measure on  $(\Gamma, \mathcal{G})$ . The set  $\Gamma$  is  $(n, \varepsilon, \kappa)$ -small if for all  $z \in \Gamma$  and  $A \in \mathcal{G}$  we have that  $P^n(z, A) \geq \varepsilon \cdot \kappa(A)$ .

Indeed, we can set  $n = k + \ell$  and  $\varepsilon = \alpha + \beta$  and we get the condition of the definition.

**Corollary 1.** There is  $n \in \mathbb{N}$ ,  $\varepsilon > 0$ , and  $\kappa$  such that  $\Gamma$  is  $(n, \varepsilon, \kappa)$ -small.

From the fact that the whole state space of a Markov chain is small, we get the desired statement using standard results on Markov chains on general state space. We get that  $\Phi$  has a unique invariant measure  $\pi$  and that the SLLN holds for  $\Phi$ , see [9, Theorem 3.6].

From the SLLN, we directly get that  $\mathbf{d} = E_\pi[\delta]$ . Now we show that  $\mathbf{c} = \frac{E_\pi[\tau]}{E_\pi[W]}$ . Let us consider a run  $(s_0, \nu_0) (s_1, \nu_1) \dots$ . By  $t_i$  we denote  $\nu_{i+1}(\Delta) - \nu_i(\Delta)$  – the time spent in the  $i$ -th state. We have

$$\begin{aligned} \mathbf{c}(\sigma) &= \lim_{n \rightarrow \infty} \frac{\sum_{i=0}^n \delta(s_i) \cdot t_i}{\sum_{i=0}^n t_i} = \lim_{n \rightarrow \infty} \frac{\sum_{i=0}^n \delta(s_i) \cdot t_i}{n} \cdot \frac{n}{\sum_{i=0}^n t_i} = \frac{\lim_{n \rightarrow \infty} (\sum_{i=0}^n \delta(s_i) \cdot t_i) / n}{\lim_{n \rightarrow \infty} (\sum_{i=0}^n t_i) / n} \\ &= \frac{E_\pi[\tau]}{E_\pi[W]} \end{aligned}$$

The fact that  $\mathbf{c}(\sigma)$  is well-defined follows from the end which justifies the manipulations with the limits. It remains to explain the last equality. First, let us divide the space of configurations into a grid  $C_\delta$ . Each  $\square \in C_\delta$  is a hypercube of configurations of unit length  $\delta$ . By  $z_i$ , we denote the  $i$ -th configuration of the run. We obtain

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{\sum_{i=0}^n t_i}{n} &= \lim_{n \rightarrow \infty} \sum_{\square \in C_\delta} \frac{\sum_{i=0}^n \mathbf{1}[z_i \in \square] \cdot t_i}{n} \\ &= \sum_{\square \in C_\delta} \lim_{n \rightarrow \infty} \frac{\sum_{i=0}^n \mathbf{1}[z_i \in \square] \cdot t_i}{\sum_{i=0}^n \mathbf{1}[z_i \in \square]} \cdot \lim_{n \rightarrow \infty} \frac{\sum_{i=0}^n \mathbf{1}[z_i \in \square]}{n} = (*) \end{aligned}$$

The second limit equals by the SLLN to  $\pi(\square)$ . By taking  $\delta \rightarrow 0$  we get that  $(*) = E_\pi[W]$ .

By similar arguments we also get that

$$\lim_{n \rightarrow \infty} \frac{\sum_{i=0}^n \delta(s_i) \cdot t_i}{n} = E_\pi[\tau]$$

□

For the proof of Lemma 3 we introduce several definitions and two auxiliary lemmata.

**Definition 6.** A path  $(s_0, v_0) \cdots (s_n, v_n)$  is  $\delta$ -wide if for every  $0 \leq i \leq n$  the configuration  $(s_i, v_i)$  is  $\delta$ -separated and for every  $0 \leq i < n$  any every bounded variable-delay event  $e \in \mathbf{E}(s_i)$  we have  $v_i(e) + v_{i+1}(\Delta) < u_e - \delta$ , i.e. no variable-delay event gets  $\delta$  close to its upper bound.

We say that a path  $(s_0, v_0) \cdots (s_n, v_n)$  has a trace  $\bar{s}_0 E_1 \bar{s}_1 E_1 \cdots E_n \bar{s}_n$  if  $\bar{s}_i = s_i$  for every  $0 \leq i \leq n$  and for every  $0 < i \leq n$  we can get from  $(s_{i-1}, v_{i-1})$  to  $(s_i, v_i)$  via occurrence of the set of events  $E_i$  after time  $v_i(\Delta)$ .

A path  $(s_0, v_0) \cdots (s_n, v_n)$  has a total time  $t$  if  $t = \sum_{i=1}^n v_i(\Delta)$ .

The idea is that a  $\delta$ -wide path can be approximately followed with positive probability. Furthermore, as formalized by the next lemma, if we have different  $\delta$ -wide paths to the same configuration  $z^*$  that have the same length and the same trace, we have similar  $n$ -step behavior (on a set of states specified by some measure  $\kappa$ ).

**Lemma 7.** For any  $\delta > 0$ , any  $n \in \mathbb{N}$ , any configuration  $(s_n, v_n)$ , and any trace  $s_0 E_1 \cdots E_n s_n$  there is a probability measure  $\kappa$  and  $\beta > 0$  such that the following holds. For every  $\delta$ -wide path  $(s_0, v_0) \cdots (s_n, v_n)$  with trace  $T = s_0 E_1 \cdots E_n s_n$  and total time  $t \geq M$  and for every  $Y \in \mathcal{G}$  we have  $P^n((s_0, v_0), Y) \geq \beta \cdot \kappa(Y)$ .

*Proof.* Recall that  $B = \max(\{\ell_e, u_e \mid e \in \mathcal{E}\} \setminus \infty)$ . Notice that the assumptions on the events' densities imply that all delays' densities are bounded by some  $c > 0$  in the following sense. For every  $e \in \mathcal{E}$  and for all  $x \in [0, B]$ ,  $d(x) > c$  or equals 0. Similarly,  $\int_B^\infty d(x)dx > c$  or equals 0.

We will find a set of configurations  $Z$  “around” the state  $z_n = (s_n, v_n)$  and define the probability measure  $\kappa$  on this set  $Z$  such that  $\kappa(Z) = 1$ . Then we show for each measurable  $Y \subseteq Z$  the desired property.

Intuitively, configurations around  $z_n$  are of the form  $(s_n, v')$  where each  $v'(e)$  is either exactly  $v(e)$  or in a small interval around  $v_n(e)$ . We now discuss which case applies to which event  $e$  for a fixed trace  $T$ . All the following notions are defined with respect to  $T$ . We say that the ticking event  $g$  is *active until the  $i$ -th step* if  $g \in \mathbf{E}(s_0) \cap \cdots \cap \mathbf{E}(s_{i-1})$ . We say that an event  $e \in \mathbf{E}(s_n) \cup E_n \cup \{\Delta\}$  is *originally scheduled in the  $i$ -th step* by  $f$  if

- either  $f = g$  and  $g$  is active until the  $i$ -th step or  $f$  is a variable-delay event; and
- there is  $k \geq 1$  and a chain of events  $e_1 \in E_{c_1}, \dots, e_k \in E_{c_k}$  such that  $e_1 = f$ ,  $c_1 = i$ , all  $e_2, \dots, e_k$  are fixed-delay events, occurrence of each  $E_{c_i}$  newly schedules  $e_{i+1}$ , occurrence of  $E_{c_k}$  newly schedules  $e$ , and  $e \in \mathbf{E}(s_{c_k}) \cap \cdots \cap \mathbf{E}(s_{n-2}) \cap (\mathbf{E}(s_{n-1}) \cup \{\Delta\})$ .

Recall that the special valuation symbol  $\Delta$  denoting the length of the last step is also part of the state space. Notice that in the previous definition, we treat  $\Delta$  as an event that is scheduled only in the state  $s_{n-1}$ . We say that *the last step is variable* if  $E_n$  is either a

singleton of a variable-delay event or all the events in  $E_n$  are originally scheduled by a variable-delay event. Otherwise, we say that *the last step is fixed*.

Intuitively, we cannot alter the value of an event  $e$  on the trace  $T$  (i.e.,  $v'(e) = v(e)$ ) if the last step is fixed and  $e$  is originally scheduled by the ticking event. In all other cases, the value of  $e$  can be altered such that  $v'(e)$  lies in a small interval around  $v_n(e)$ . The rest of the proof is divided in two cases.

**The last step is fixed** Let us divide the events  $e \in \mathbf{E}(s_n) \cup \{\Delta\}$  into three sets as follows

- $e \in A$  if  $e$  is originally scheduled by a variable-delay event and  $\text{frac}(v_n(e)) \neq 0$ ;
- $e \in B$  if  $e$  is originally scheduled by a variable-delay event and  $\text{frac}(v_n(e)) = 0$ ;
- $e \in C$  if  $e$  is originally scheduled by the ticking event.


Let  $a_1, \dots, a_d$  be the distinct fractional values of the events  $A$  in the valuation  $v_n$  ordered increasingly by the step in which the corresponding events were originally scheduled. This definition is correct because two events with the same fractional value must be originally scheduled by the same event in the same step. Furthermore, let  $F_1, \dots, F_d$  be the corresponding sets of events, i.e.  $\text{frac}(v_n(e_i)) = a_i$  for any  $e_i \in F_i$ . We call a configuration  $z \sim z_n$  such that all events  $e \in (B \cup C)$  have the same value in  $z$  and  $z_n$  a *target* configuration and treat it as a  $d$ -dimensional vector describing the distinct values for the sets  $F_1, \dots, F_d$ . A  $\delta$ -neighborhood of a target configuration  $z$  is the set of configuration  $\{z + C \mid C \in (-\delta, \delta)^d\}$ . Observe that the  $\delta$ -neighborhood is a  $d$ -dimensional space. We set  $Z$  to be the  $(\delta/4)$ -neighborhood of  $z_n$  (the reason for dividing  $\delta$  by 4 is technical and will become clear in the course of this proof). Let  $\kappa_d$  denote the standard Lebesgue measure on the  $d$ -dimensional affine space and set  $\kappa(Y) := \kappa_d(Y)/\kappa_d(Z)$  for any any measurable  $Y \subseteq Z$ .

In order to prove the probability bound for any measurable  $Y \subseteq Z$ , it suffices to prove it for the generators of  $Z$ , i.e. for  $d$ -dimensional hypercubes centered around some state in  $Z$ . Let us fix an arbitrary  $z \in Z$  and  $\gamma < \delta/4$ . We set  $Y$  to be the  $\gamma$ -neighborhood of  $z$ . In the rest of the proof we will show how to reach the set  $Y$  from the initial state  $(s_0, v_0)$  in  $n$  steps with high enough probability.

We show it by altering the original  $\delta$ -wide path  $\sigma = (s_0, v_0) \cdots (s_n, v_n)$ . Let  $t_1, \dots, t_n$  be the waiting times such that  $t_i = v_i(\Delta)$ . In the first phase, we reach the fixed  $z$  instead of the configuration  $z_n$ . We find waiting times  $t'_1, \dots, t'_n$  that induce a path  $\sigma' = (s_0, v_0) (s_1, v'_1) \dots (s_n, v'_n)$  with trace  $T$  such that  $(s_n, v'_n) = z$  and  $t'_i = v'_i(\Delta)$ . In the second phase, we define using  $\sigma'$  a set of paths to  $Y$ . We allow for intervals  $I_1, \dots, I_n$  such that for any choice  $\tilde{t}_1 \in I_1, \dots, \tilde{t}_n \in I_n$  we get a path  $\tilde{\sigma} = (s_0, v_0) (s_1, \tilde{v}_1) \dots (s_n, \tilde{v}_n)$  such that  $(s_n, \tilde{v}_n) \in Y$  and  $\tilde{t}_i = \tilde{v}_i(\Delta)$ . From the size of the intervals for variable-delay events and from the bound on densities  $c$  we get the overall bound on probabilities. Let us start with the first step.

Let  $v_1, \dots, v_d$  be the distinct values of the target configuration  $z$ . Recall that  $|v_i - a_i| < \delta/4$  for each  $i$ . Let  $r(1), \dots, r(d)$  be the indices such that all events in  $F_i$  are originally scheduled in the step  $r(i)$ . Notice that each  $E_{r(i)}$  is a singleton of a variable-delay event.






**Fig. 4.** Illustration of paths leading to the set  $Y$ . The original path  $\sigma$  is in the first phase altered to reach the target state  $z$  (its values  $v_1, v_2$ , and  $v_3$  are depicted between  $\sigma$  and  $\sigma'$ ). In the second phase, a set of paths that reach  $Y$  is constructed by allowing imprecision in the waiting times – the transition times are randomly chosen inside the hatched areas. Notice that at most  $d$  smaller intervals of size  $\gamma/2$  can be used to get constant probability bound with respect to the size of the  $d$ -dimensional hypercube  $Y$ . Transitions with fixed-delay are omitted from the illustration (except for the last transition).

As illustrated in Figure 4, we set for each  $1 \leq i \leq m$

$$t'_i = \begin{cases} \ell_e - v_{i-1}(e) & \text{if } e \in E_i \text{ is fixed-delay,} \\ t_i + \sum_{k=1}^{i-1} (t_k - t'_k) + a_j - v_j & \text{if } i = r(j) \text{ for } 1 \leq j \leq d, \\ t_i + \sum_{k=1}^{i-1} (t_k - t'_k) & \text{otherwise.} \end{cases}$$

Intuitively, we adjust the variable-delays in the steps preceding the original scheduling of sets  $F_1, \dots, F_d$  whereas the remaining variable-delay steps are kept in sync with the original path  $\sigma$ . The absolute time of any transition in  $\sigma'$  (i.e. the position of a line depicting a configuration in Figure 4) is not shifted by more than  $\delta/4$  since  $|v_i - a_i| < \delta/4$  for any  $i$ . Thus, the difference of any two absolute times is not changed by more than  $\delta/2$ . This difference bounds the difference of  $|v_i(e) - v'_i(e)|$  for any  $i$  and  $e \in \mathcal{E}$ . Hence,  $\sigma'$  is  $(\delta/2)$ -wide because  $\sigma$  is  $\delta$ -wide. Furthermore,  $\sigma'$  goes through the same regions as  $\sigma$  and performs the same sequence of events scheduling. Building on that, the desired property  $z'_n = z$  is easy to show.

Next we allow imprecision in the waiting times of  $\sigma'$  so that we get a set of paths of measure linear in  $\gamma^d$ . In each step we compensate for the imprecision of the previous



**Fig. 5.** Illustration of construction of  $\bar{\sigma}$  for the empty set  $C$  and the last step variable.

step. Formally, let  $T_i$  denote  $t'_i + \sum_{k=1}^{i-1} (t'_k - \bar{t}_k)$ . For each  $1 \leq i \leq m$  we constraint

$$\bar{t}_i \in \begin{cases} [T_i, T_i] & \text{if } E_i \text{ are fixed-delay events,} \\ (T_i - \frac{\gamma}{2}, T_i + \frac{\gamma}{2}) & \text{if } i = r(j) \text{ for } 1 \leq j \leq d, \\ (T_i - \frac{\delta}{4}, T_i + \frac{\delta}{4}) & \text{otherwise.} \end{cases}$$

The difference to  $\sigma'$  of any two absolute times is not changed by more than  $\delta/2$  because the imprecision of any step is bounded by  $\delta/4$ . Because  $\sigma'$  is  $(\delta/2)$ -wide, any path  $\bar{\sigma}$  goes through the same regions as  $\sigma'$ . The difference of the value of events in any  $F_i$  in the state  $\bar{z}_n$  from the state  $z$  is at most  $\gamma/2$  because it is only influenced by the imprecision of the step preceding its original scheduling. Hence,  $\bar{z}_n \in Y$ .

By  $v$  we denote the number of variable-delay singletons among  $E_1, \dots, E_n$ . From the definition of  $P$ , it is easy to prove by that

$$P^n(z_0, Y) \geq p_{\min}^n \cdot (c \cdot \gamma)^d \cdot (c \cdot \delta/2)^{v-d} \geq (p_{\min} \cdot c/2)^n \cdot \gamma^d \cdot \delta^{n-d}$$

Since  $\kappa_d(Y) = (2 \cdot \gamma)^d$  and  $\kappa_d(Z) = (2 \cdot \delta/4)^d$ , we have  $\kappa(Y) = \kappa_d(Y)/\kappa_d(Z) = (4\gamma/\delta)^d$ . We get  $P^n(z_0, Y) \geq \kappa(Y) \cdot (\delta \cdot p_{\min} \cdot c/8)^n$  and conclude the proof of this case by setting  $\varepsilon = (\delta \cdot p_{\min} \cdot c/8)^n$ .

**The last step is variable** The rest of the proof proceeds in a similar fashion as previously, we reuse the same notions and the same notation. We only redefine the differences: the neighbourhood and the way the paths are altered.

We call  $(s, v) \sim z_n$  a *target* configuration if there is  $y \in \mathbb{R}$  such that for all events  $e \in C$  we have  $v(e) - v_n(e) = y$  and for all events  $e \in B$  we have  $v(e) = v_n(e)$ . We set  $g = d + 1$  if  $C$  is non-empty, and  $g = d$ , otherwise. We treat a target configuration as a  $g$ -dimensional vector describing the distinct values for the sets  $F_1, \dots, F_d$  and the value  $y$ , if necessary. Again, a  $\delta$ -neighborhood of a target configuration  $z$  is the set of configuration  $\{z + C \mid C \in (-\delta, \delta)^g\}$ . We set  $Z$  to be the  $(\delta/4)$ -neighborhood of  $z_n$  and set  $\kappa(Y) := \kappa_g(Y)/\kappa_g(Z)$  for any any measurable  $Y \subseteq Z$ . We fix  $Y$  to be a  $\gamma$ -neighborhood of a fixed  $z \in Z$ .

The path  $\sigma'$  is obtained from the  $\sigma$  in the same way as before. We need to allow imprecision in the waiting times of  $\sigma'$  so that we get a set of paths of measure linear in  $\gamma^g$ .

- For the case  $g = d + 1$  it is straightforward as we make the last step also with imprecision  $\pm\gamma/2$ . Precisely

$$\bar{t}_i \in \begin{cases} [T_i, T_i] & \text{if } E_i \text{ are fixed-delay events,} \\ (T_i - \frac{\gamma}{2}, T_i + \frac{\gamma}{2}) & \text{if } i = r(j) \text{ for } 1 \leq j \leq d \text{ or } i = m, \\ (T_i - \frac{\delta}{4}, T_i + \frac{\delta}{4}) & \text{otherwise} \end{cases}$$

where  $E_n$  are originally scheduled in the  $m$ -th step if  $E_n$  are fixed-delay events, and  $m$  equals  $n$ , otherwise. The difference of the value of events in any  $F_i$  in the state  $\bar{z}_n$  from the state  $z$  is at most  $\gamma$  because it is influenced by the imprecision of the step preceding its original scheduling and also by the imprecision of the last step. Events in  $C$  have the difference of the value at most  $\gamma/2$  because of the last step. Hence,  $\bar{z}_n \in Y$ . Again, we get that  $P^n(z_0, Y) \geq \kappa(Y) \cdot (\delta \cdot p_{\min} \cdot c/8)^n$  and conclude the proof by setting  $\varepsilon = (\delta \cdot p_{\min} \cdot c/8)^n$ .

- For the case  $g = d$  it is somewhat tricky since only at most  $d$  choices of waiting times can have their precision dependent on  $\gamma$ . In each step we compensate for the imprecision of the previous step. Only the imprecision of the step preceding the first scheduling  $E_1$  is not compensated for. Otherwise, it would influence the value of events  $E_1$  in  $\bar{z}_n$ . Let  $T_i^a$  denote  $t'_i + \sum_{k=a}^{i-1} (t'_k - \bar{t}_k)$ . As illustrated in Figure 5, we constraint

$$\bar{t}_i \in \begin{cases} [T_i^1, T_i^1] & \text{if } E_i \text{ are fixed-delay events,} \\ (T_i^1 - \frac{\delta}{6}, T_i^1 + \frac{\delta}{6}) & \text{if } i \leq r(1), \\ (T_i^{r(1)+1} - \frac{\gamma}{2}, T_i^{r(1)+1} + \frac{\gamma}{2}) & \text{if } i = r(j) \text{ for } 2 \leq j \leq d \text{ or } i = m, \\ (T_i^{r(1)+1} - \frac{\delta}{6}, T_i^{r(1)+1} + \frac{\delta}{6}) & \text{otherwise.} \end{cases}$$

The difference to  $\sigma'$  of any two absolute times is not changed by more than  $3 \cdot \delta/6 = \delta/2$  because the imprecision of any step is bounded by  $\delta/6$ . Because  $\sigma'$  is  $(\delta/2)$ -wide, any path  $\bar{\sigma}$  goes through the same regions as  $\sigma'$ . The difference of the value of events  $E_1$  in the state  $\bar{z}_n$  from the state  $z$  is at most  $\gamma/2$  because it is only influenced by the imprecision of the last step. The difference of any other event  $e$  is at most  $2 \cdot \gamma/2$  because it is influenced by the imprecision of the step preceding the original scheduling of  $e$ , as well. Hence,  $\bar{z}_n \in Y$ .

Now, we get that  $P^n(z_0, Y) \geq \kappa(Y) \cdot (\delta \cdot p_{\min} \cdot c/12)^n$  and conclude the proof by setting  $\varepsilon = (\delta \cdot p_{\min} \cdot c/12)^n$ .  $\square$

**Lemma 8.** *Let  $\delta > 0$  and  $R$  be a region such that the ticking event  $e$  is either not scheduled or has the greatest value among all events scheduled in  $R$ . There is  $n \in \mathbb{N}$ ,  $\delta' > 0$ , a configuration  $z^*$ , and a trace  $s_0 E_1 \cdots E_n s_n$  such that from any  $\delta$ -separated  $z \in R$ , there is a  $\delta'$ -wide path to  $z^*$  with trace  $s_0 E_1 \cdots E_n s_n$  and total time  $t \geq M$ .*

*Proof.* We use a similar concept as in the proof of Lemma 1. Let us fix a  $\delta$ -separated  $z \in R$ . Let  $a$  be the greatest value of all event scheduled in  $z$ . Observe, that no value is in the interval  $(a, a + \delta)$ . When we build the  $\delta$ -wide path step by step, we use a variable  $s$  denoting start of this interval of interest which flows with time. Before the first step, we have  $s := a$ . After each step, which takes  $t$  time, we set  $s := \text{frac}(s + t)$ .

In the interval  $[s, s + \delta]$  we make a grid of  $3 \cdot |\mathcal{E}| + 1$  points that we shift along with  $s$ , and set  $\delta' = \delta / (3 \cdot |\mathcal{E}| + 1)$ . On this grid, a procedure similar to the  $\delta$ -separation takes place. We build the  $\delta'$ -wide path by choosing sets of events  $E_i$  to occur, waiting times  $t_i$  of the individual transitions, and target states  $z_i$  after each transition so that

- every variable-delay event occurs exactly at an empty point of the grid (i.e. at a time when an empty point has value 0), and
- the built path is “feasible”, i.e. all the specified events can occur after the specified waiting time, and upon each occurrence of a specified event we move to the specified target state with positive probability,

These rules guarantee that the path we create is  $\delta'$ -wide. Indeed, the initial configuration is  $\delta$ -separated for  $\delta > \delta'$ , upon every new transition, the  $\delta'$ -neighborhood of 0 is empty, and every variable-delay event occurs at a point different from its current point, whence it occurs at least  $\delta'$  prior to its upper bound. It is easy to see that such choices are possible since there are only  $\mathcal{E}$  events, but  $3 \cdot |\mathcal{E}| + 1$  points.

Now we show that this procedure lasts only a fixed amount of steps before all the scheduled events lie on the grid. Notice that if the ticking event is scheduled in  $R$ , it lies at a point of the grid from the very beginning because we define the grid adjacent to its value. If it is not scheduled, it can get scheduled only by a variable-delay event which occurs already at a point of the grid. Values of any other scheduled fixed-delay event gets eventually placed at a point of a grid. Indeed, every such event gets scheduled by a variable-delay event next time, since we assume a single-ticking GSMP. We now that after time  $M$ , all the non-ticking fixed-delay events are either not scheduled or lie on the grid. Each step takes at least  $\delta'$  time. In total, after  $n = \lceil M/\delta' \rceil + 1$  steps with trace  $E_1, \dots, E_n$ , we can set  $z^* := z_n$ .

It remains to show that from any other  $\delta'$ -separated configuration  $z' \in R$ , we can build a  $\delta'$ -wide path of length  $n$ , with trace  $E_1, \dots, E_n$  that ends in  $z^*$ . We start in the same region. From the definition of the region relation and from the fact that all events occur in the empty interval  $(a, a + \delta)$  we get the following. By appropriately adjusting the waiting times so that the events occur at the same points of the grid as before, we can follow the same trace and the same control states (going through the same regions) and build a path  $z'_0 \dots z'_n$  such that  $z'_n = z_n$ . Indeed, all scheduled events have the same value in  $z'_n$  as in  $z_n$  because they lie on the same points of the grid. In fact, this holds for  $z'_{n-1}$  and  $z_{n-1}$  as well (because the first  $n - 1$  steps take more than  $M$  time units) except for the value of  $\Delta$ . Finally, also  $\Delta$  has the same value in  $z'_n$  as in  $z_n$  because there is no need to alter the waiting time in the last step. By the same arguments as before, the built path is also  $\delta'$ -wide.  $\square$

**Lemma 3.** There is a measurable set of configurations  $C$  such that

1. there is  $k \in \mathbb{N}$  and  $\alpha > 0$  such that for every  $z \in F$  we have  $P^k(z, C) \geq \alpha$ , and

2. there is  $\ell \in \mathbb{N}$ ,  $\beta > 0$ , and a probability measure  $\kappa$  such that for every  $z \in C$  and  $A \in \mathcal{G}$  we have  $P^\ell(z, A) \geq \beta \cdot \kappa(A)$ .

*Proof.* We choose some reachable region  $R$  such that the ticking event  $e$  is either not scheduled in  $R$  or  $e$  has the greatest fractional part among all the scheduled events. There clearly is such a region. We fix a sufficiently small  $\delta > 0$  and choose  $C$  to be the set of  $\delta$ -separated configurations in  $R$ . Now, we show how we fix this  $\delta$ .

It is a standard result from the theory of Markov chains, see e.g. [25, Lemma 8.3.9], that in every ergodic Markov chain there is  $n$  such that between any two states there is a path of length exactly  $n$ . The same result holds for the aperiodic region graph  $G$ . From Lemma 1, we reach in  $m$  steps from any  $z \in \Gamma$  a  $\delta'$ -separated configuration  $z'$  with probability at least  $q$ . From  $z'$ , we have a path of length  $n$  to the region  $R$ . From Lemma 2, we have  $p > 0$  such that we reach  $R$  from  $z'$  in  $n$  steps with probability at least  $p$ . Furthermore, we end up in a  $(\delta'/3^n)$ -separated configuration of the region  $R$ . Hence, we set  $\delta = \delta'/3^n$  and obtain the first part of the lemma.

The second part of the lemma is directly by connecting Lemmata 8 and 7.

## C Proof of Theorem 5

It only remains to prove the two Claims.

*Claim.* Let  $X$  be a set of all configurations in a BSCC  $\mathcal{B}$ ,  $X_{\hat{s}} \subseteq X$  the set of configurations with state  $\hat{s}$ , and  $d$  the frequency corresponding to  $\mathcal{B}$ . There are computable constants  $n_1, n_2 \in \mathbb{N}$  and  $p_1, p_2 > 0$  such that for every  $i \in \mathbb{N}$  and  $z_X \in X$  we have

$$\begin{aligned} |\mathcal{P}(\text{Reach}(X)) - P^i(z_0, X)| &\leq (1 - p_1)^{\lfloor i/n_1 \rfloor} \\ |d - P^i(z_X, X_{\hat{s}})| &\leq (1 - p_2)^{\lfloor i/n_2 \rfloor} \end{aligned}$$

*Proof.* Let  $Y$  denote the union of regions from which the BSCC  $\mathcal{B}$  is reachable. By Lemmata 1 and 2 we have  $p, q > 0$  and  $m \in \mathbb{N}$  and  $k < |V|$  such that from any  $z \in Y$  we reach  $X$  in  $m + k$  steps with probability at least  $p \cdot q$ . We get the first part by setting  $n_1 = m + k$  and  $p_1 = p \cdot q$ . Indeed, if the process stays in  $Y$  after  $n_1$  steps, it has the same chance to reach  $X$  again, if the process reaches  $X$ , it never leaves it, and if the process reaches  $\Gamma \setminus (X \cup Y)$ , it has no chance to reach  $X$  any more.

By Corollary 1 in Appendix B.3,  $\Gamma$  is  $(n, \varepsilon, \kappa)$ -small. By Theorem 8 of [24] we thus obtain that for all  $x \in \Gamma$  and all  $i \in \mathbb{N}$ ,

$$\sup_{A \in \mathcal{G}} |P^i(x, A) - \pi(A)| \leq (1 - \varepsilon)^{\lfloor i/n \rfloor}$$

which yields the second part by setting  $A = \{(s, v) \in \Gamma \mid s = \hat{s}\}$  and observing  $d = \pi(A)$  and  $A \in \mathcal{G}$   $\square$

*Claim.* On each region,  $W$  is continuous, and  $E_\pi[W]$  is finite.

*Proof.* Let  $(s, \nu)$  be a configuration,  $C$  and  $D$  the set of variable-delay and fixed-delay events scheduled in  $s$ , respectively. For nonempty  $D$  let  $T = \min_{d \in D}(\ell_d - \nu(d))$  be the time the first fixed-delay event can occur; for  $D = \emptyset$  we set  $T = \infty$ . The probability that the transition from  $(s, \nu)$  occurs within time  $t$  is

$$F(t) = \begin{cases} 1 - \prod_{c \in C} \int_t^\infty f_{c|\nu(c)}(x) dx & \text{for } 0 < t < T, \\ 1 & \text{for } t \geq T \end{cases}$$

as non-occurrences of variable-delay events are mutually independent. Observe that  $F(t)$  is piece-wise differentiable on the interval  $(0, T)$ , we denote by  $f(t)$  its piece-wise derivative. The expected waiting time in  $(s, \nu)$  is

$$W((s, \nu)) = \begin{cases} \int_0^T t \cdot f(t) dt + T \cdot (1 - F(T)) & \text{for } T < \infty, \\ \int_0^\infty t \cdot f(t) dt & \text{for } T = \infty. \end{cases} \quad (5)$$

Recall that for each variable-delay event  $e$ , the density  $f_e$  is continuous and bounded as it is defined on a closed interval. Therefore,  $f_{e|t}$  are also continuous, hence  $F$  and  $f$  are also continuous with respect to  $\nu$  and with respect to  $t$  on  $(0, T)$ . Thus  $W$  is continuous for  $T$  both finite and infinite. Moreover, for finite  $T$ ,  $W$  is bounded by  $T$  which is for any  $(s, \nu)$  smaller than  $\max_{d \in \mathcal{E}_f} \ell_d$ . Hence,  $E_\pi[W]$  is finite. For  $T = \infty$ ,  $E_\pi[W]$  is finite due to the assumption that each  $f_e$  has finite expected value.  $\square$